



US006363129B1

(12) **United States Patent**
Agazzi

(10) **Patent No.:** **US 6,363,129 B1**
(45) **Date of Patent:** ***Mar. 26, 2002**

(54) **TIMING RECOVERY SYSTEM FOR A MULTI-PAIR GIGABIT TRANSCEIVER**

5,978,390 A * 11/1999 Balatoni 370/540
6,041,090 A * 3/2000 Chen 375/376

(75) Inventor: **Oscar E. Agazzi**, Irvine, CA (US)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **Broadcom Corporation**, Irvine, CA (US)

WO WO 98/09400 5/1998 H04L/7/02
WO WO 99/07077 2/1999

(*) Notice: This patent issued on a continued prosecution application filed under 37 CFR 1.53(d), and is subject to the twenty year patent term provisions of 35 U.S.C. 154(a)(2).

OTHER PUBLICATIONS

Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

IEEE Std 802.3ab-1999 (Supplement to IEEE Std 802.3, 1998 Edition), entitled Supplement to Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications—Physical Layer Parameters and Specifications for 1000 Mb/s Operation Over 4-Pair of Category 5 Balanced Copper Cabling, Type 1000BASE-T.

(21) Appl. No.: **09/437,721**
(22) Filed: **Nov. 9, 1999**

Search Report dated Sep. 5, 2000 relating to corresponding International Application No. PCT/US00/11123, 4 pp. Local and Metropolitan Area Networks, Specific Requirements, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications; IEEE Std 802.3, 1998 Edition; pp. 893-963: LAN MAN Standards Committee of the IEE Computer Society; USA.

Related U.S. Application Data

(60) Provisional application No. 60/130,616, filed on Apr. 22, 1999, provisional application No. 60/108,648, filed on Nov. 16, 1998, provisional application No. 60/108,319, filed on Nov. 13, 1998, and provisional application No. 60/107,874, filed on Nov. 9, 1998.

(List continued on next page.)

(51) **Int. Cl.**⁷ **H04L 7/00**
(52) **U.S. Cl.** **375/355; 375/371; 331/2; 331/46**
(58) **Field of Search** **375/216, 259, 375/260, 354, 355, 362, 371, 373, 376, 377; 370/503, 516; 331/2, 46, 48, 56, 49**

Primary Examiner—Jean Corrielus
(74) *Attorney, Agent, or Firm*—Christie, Parker & Hale, LLP

(57) **ABSTRACT**

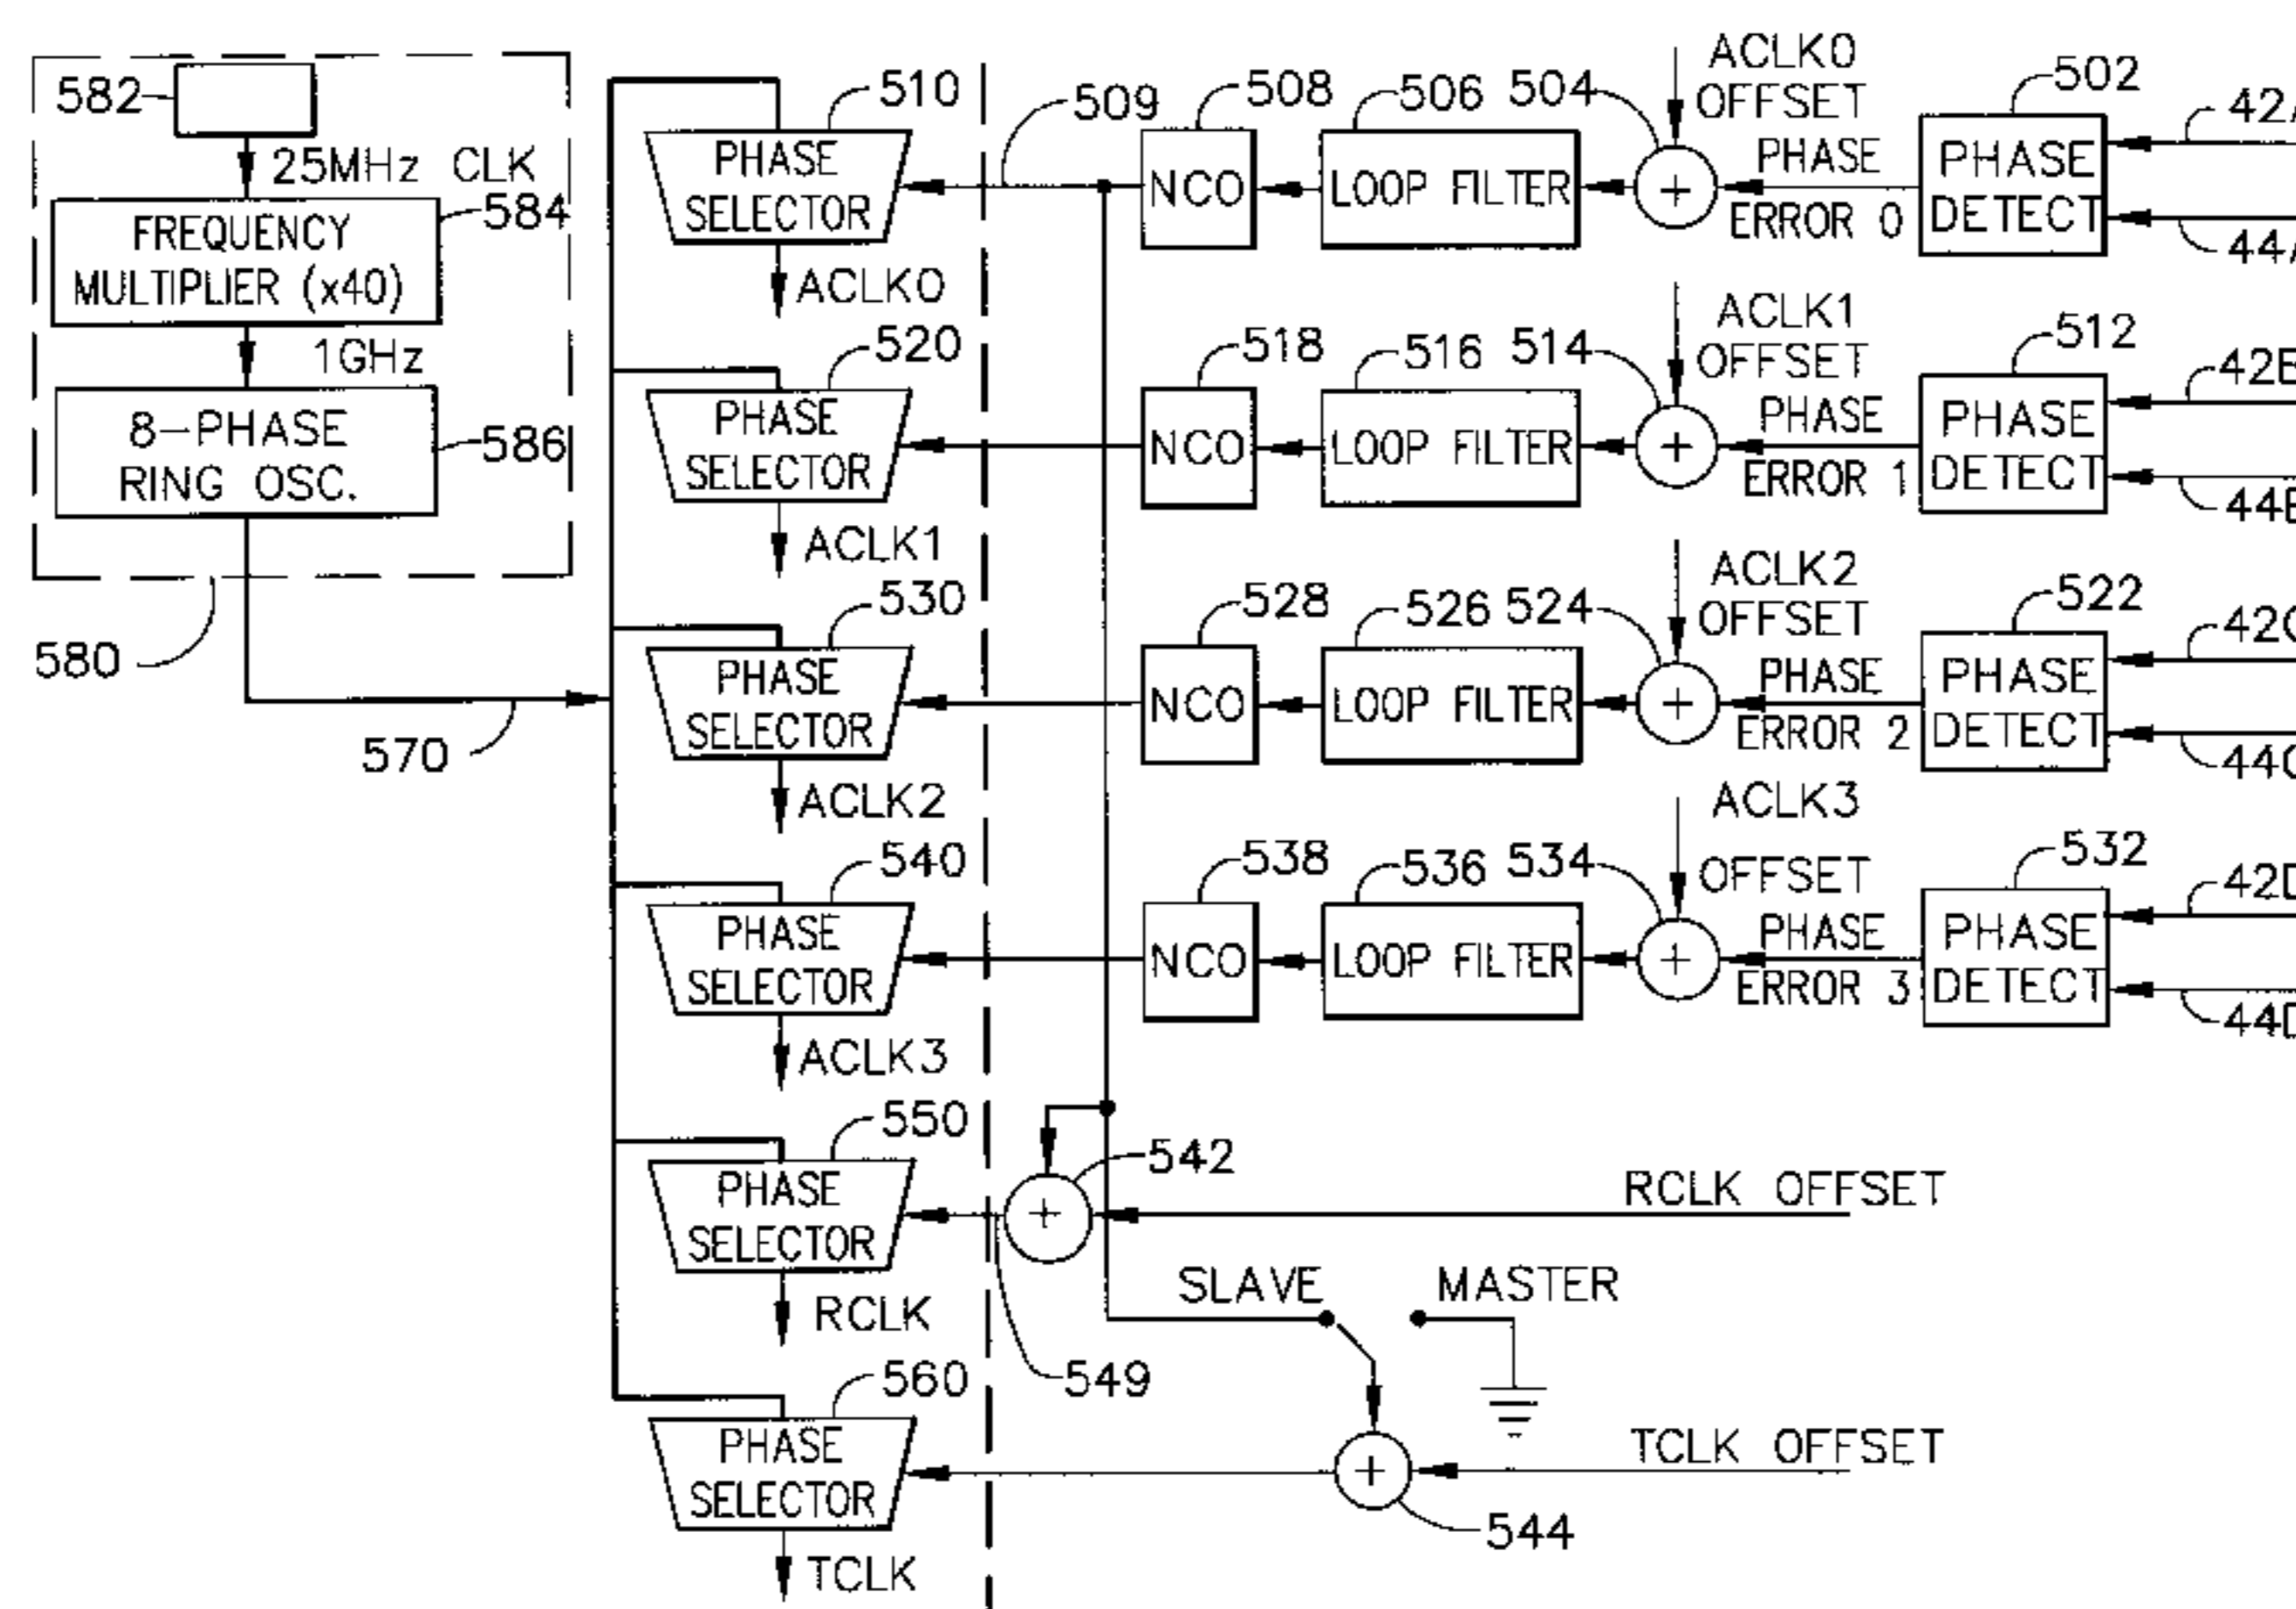
A method and a timing recovery system for generating a set of clock signals in a system which includes a set of subsystems. Each of the subsystems includes an analog section. The set of clock signals includes a set of sampling clock signals. Each of the analog sections operates in accordance with a corresponding one of the sampling clock signals. For each of the sampling clock signals, a phase error is generated from a corresponding phase detector. The phase errors are filtered by a set of corresponding loop filters. The filtered phase errors are provided to a set of corresponding oscillators to generate phase control signals. The phase control signals are provided to a set of corresponding phase selectors to generate the sampling clock signals.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,659,999 A * 4/1987 Motoyama et al.
5,128,633 A * 7/1992 Martin et al. 3313/2
5,142,377 A * 8/1992 Moriyama et al.
5,428,361 A * 6/1995 Higtower et al.
5,548,249 A * 8/1996 Sumita et al. 331/1 A
5,572,167 A * 11/1996 Alder et al. 331/2
5,644,271 A * 7/1997 Molloy et al. 331/176
5,726,607 A * 3/1998 Brede et al. 331/2

40 Claims, 13 Drawing Sheets



OTHER PUBLICATIONS

Agazzi et al.; Two-Phase Decimation and Jitter Compensation in Full-Duplex Data Transceivers; Proceedings of the International Symposium on Circuits and Systems; Mar. 10, 1992; pp. 1717-1720; IEEE; USA.

Hatamian et al.; Design Considerations for Gigabit Ethernet 1000Base-T Twisted Pair Transceivers; 1998; pp. 335-342; IEEE 1988 Custom Integrated Circuits Conference.

* cited by examiner




FIG. 1




FIG. 2




FIG. 3

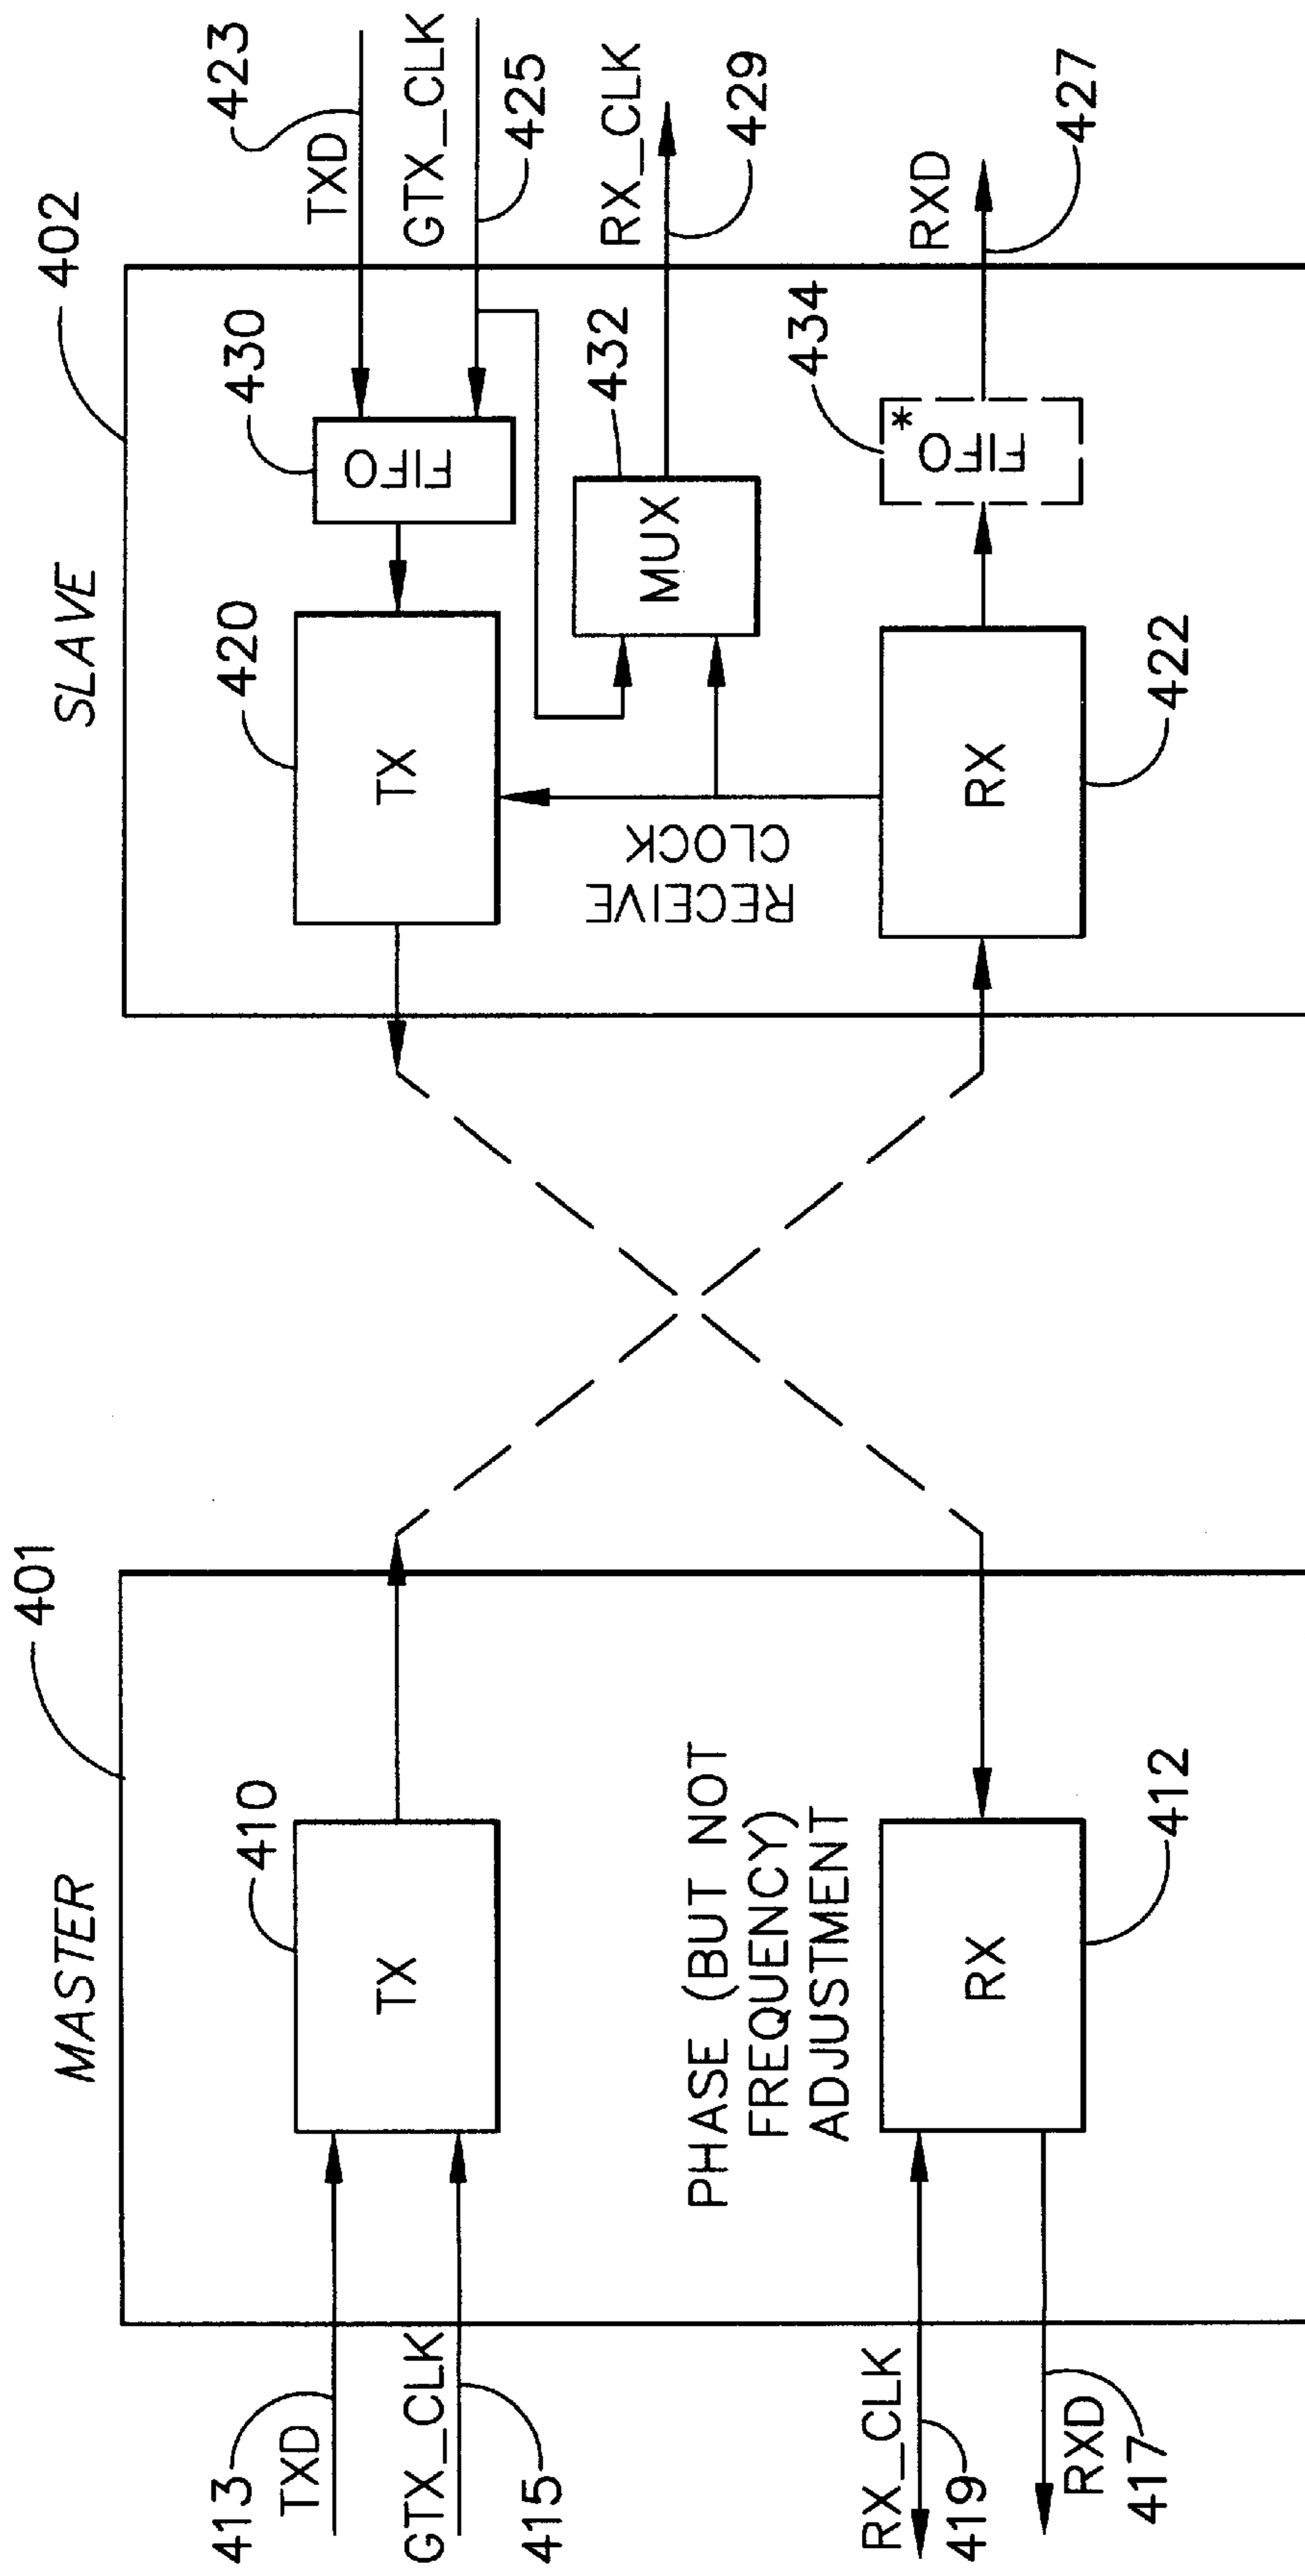


FIG. 4

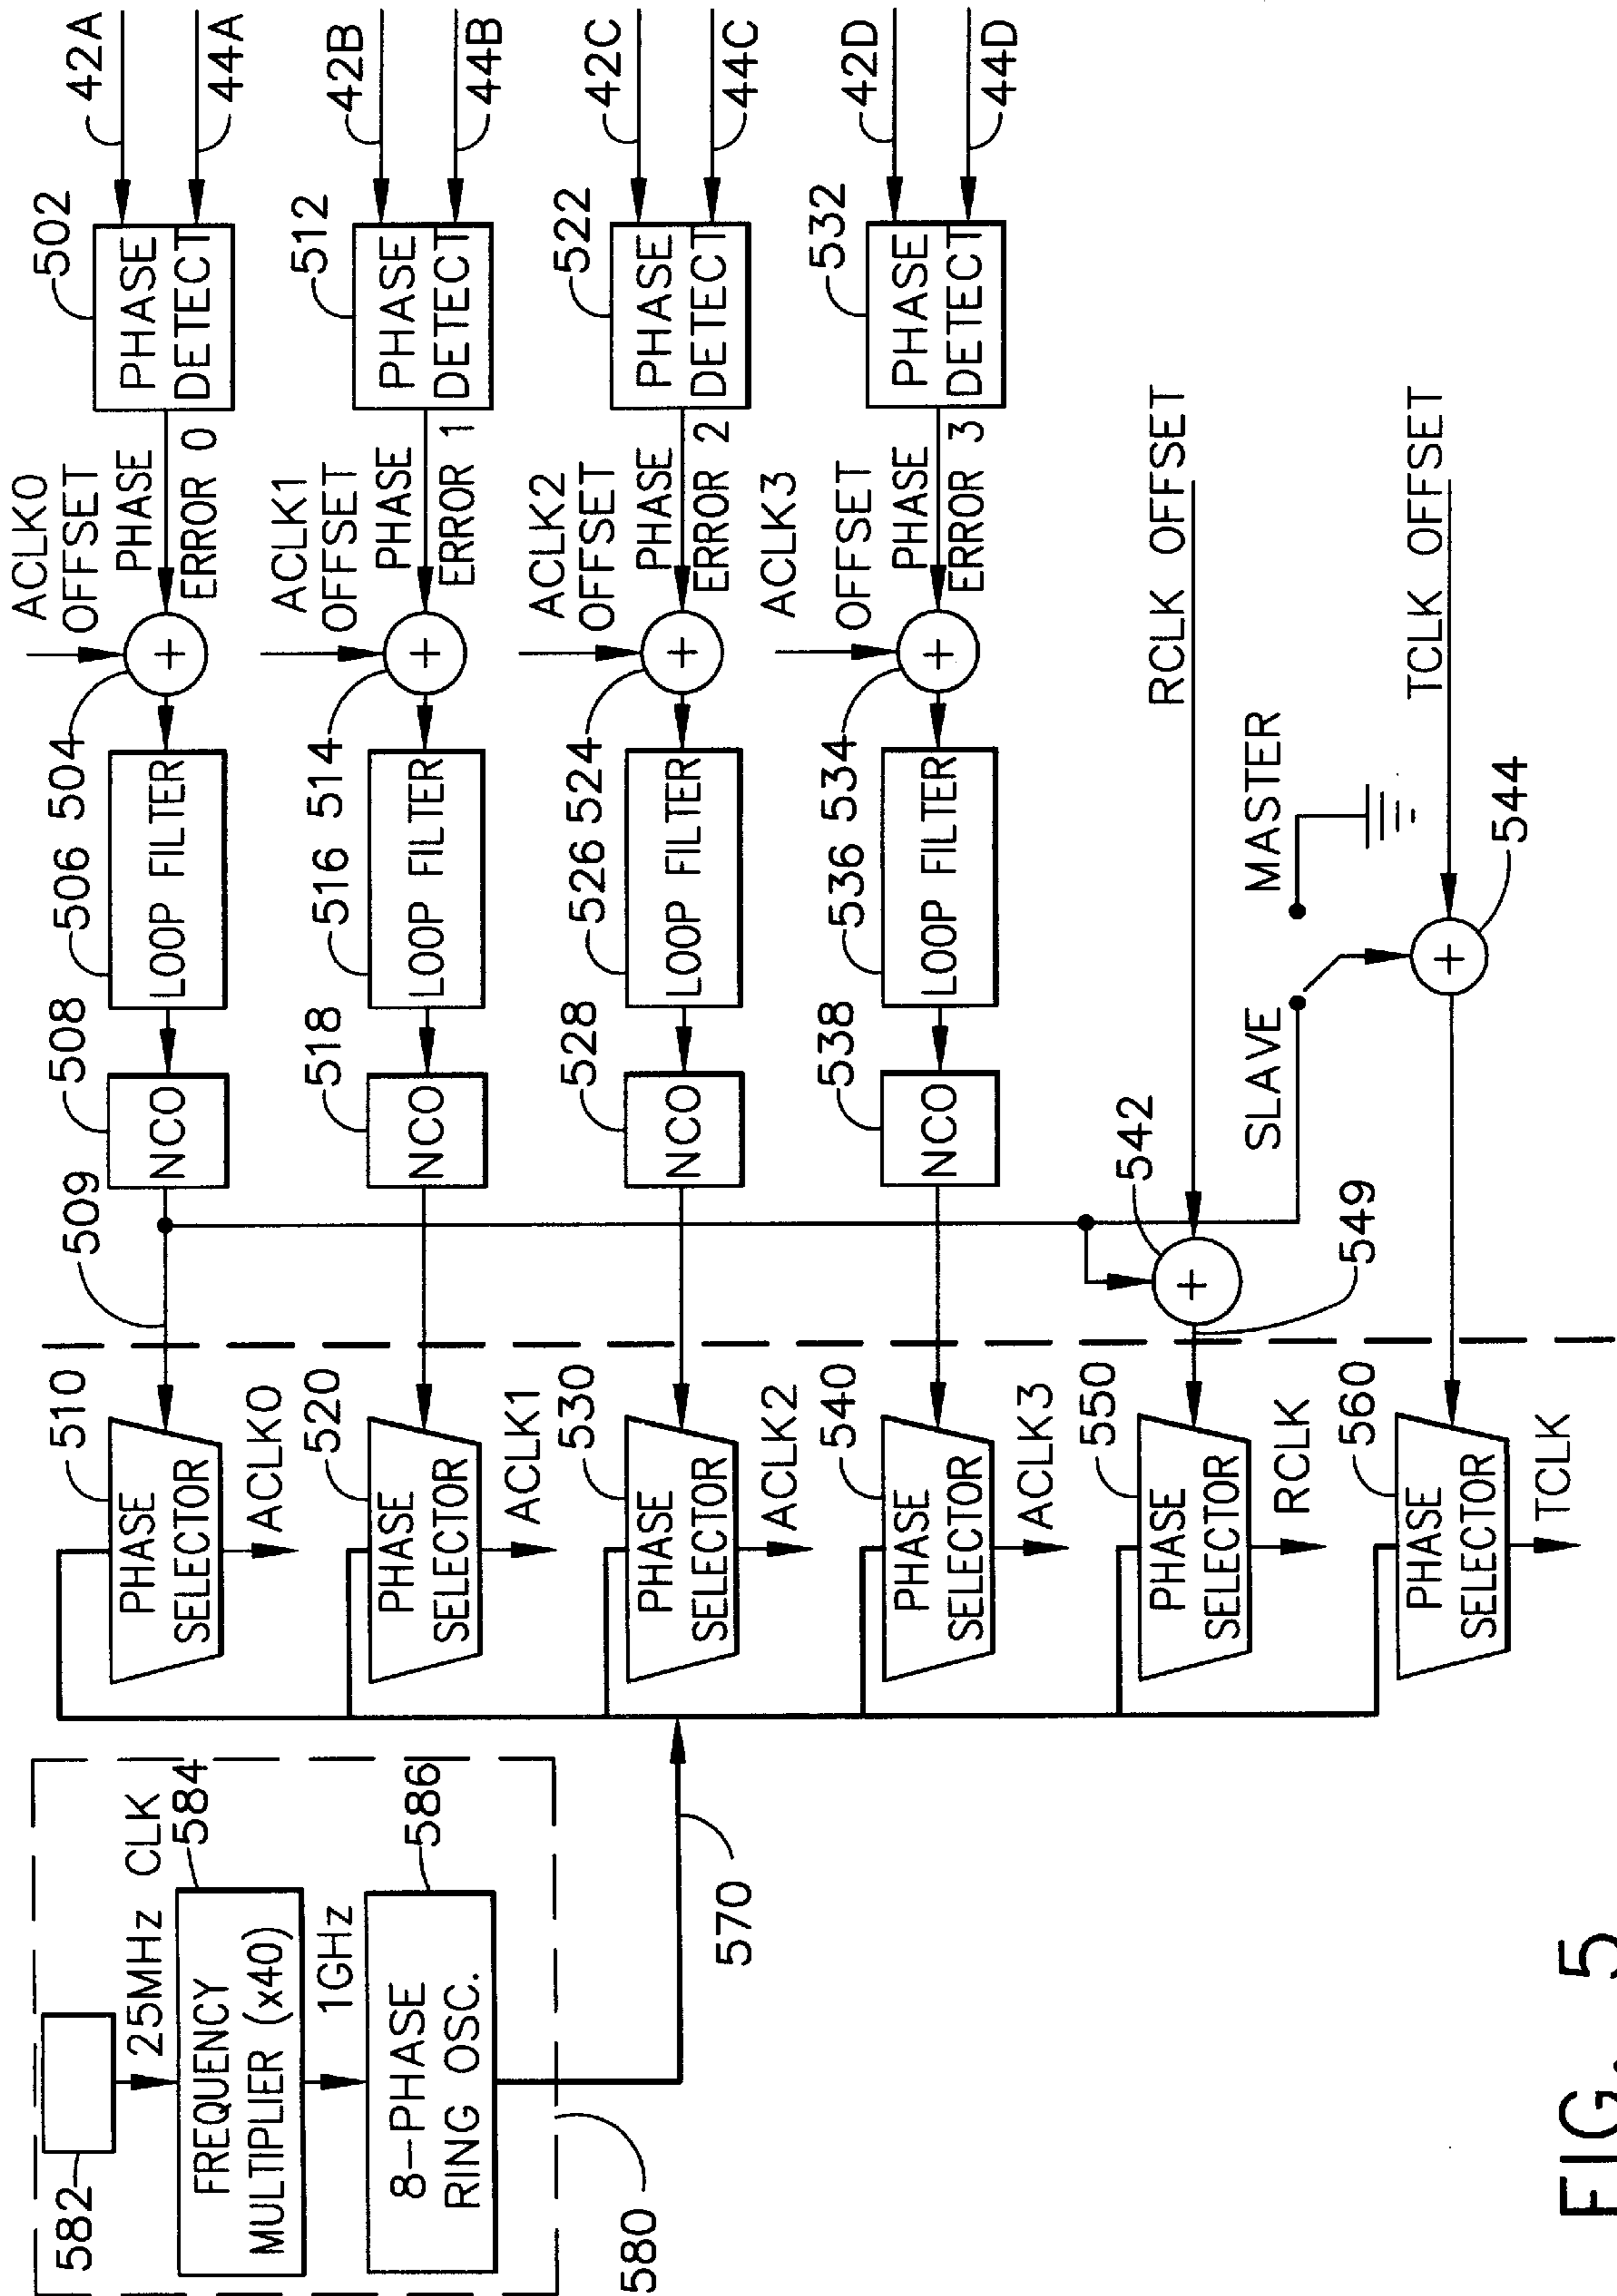


FIG. 5

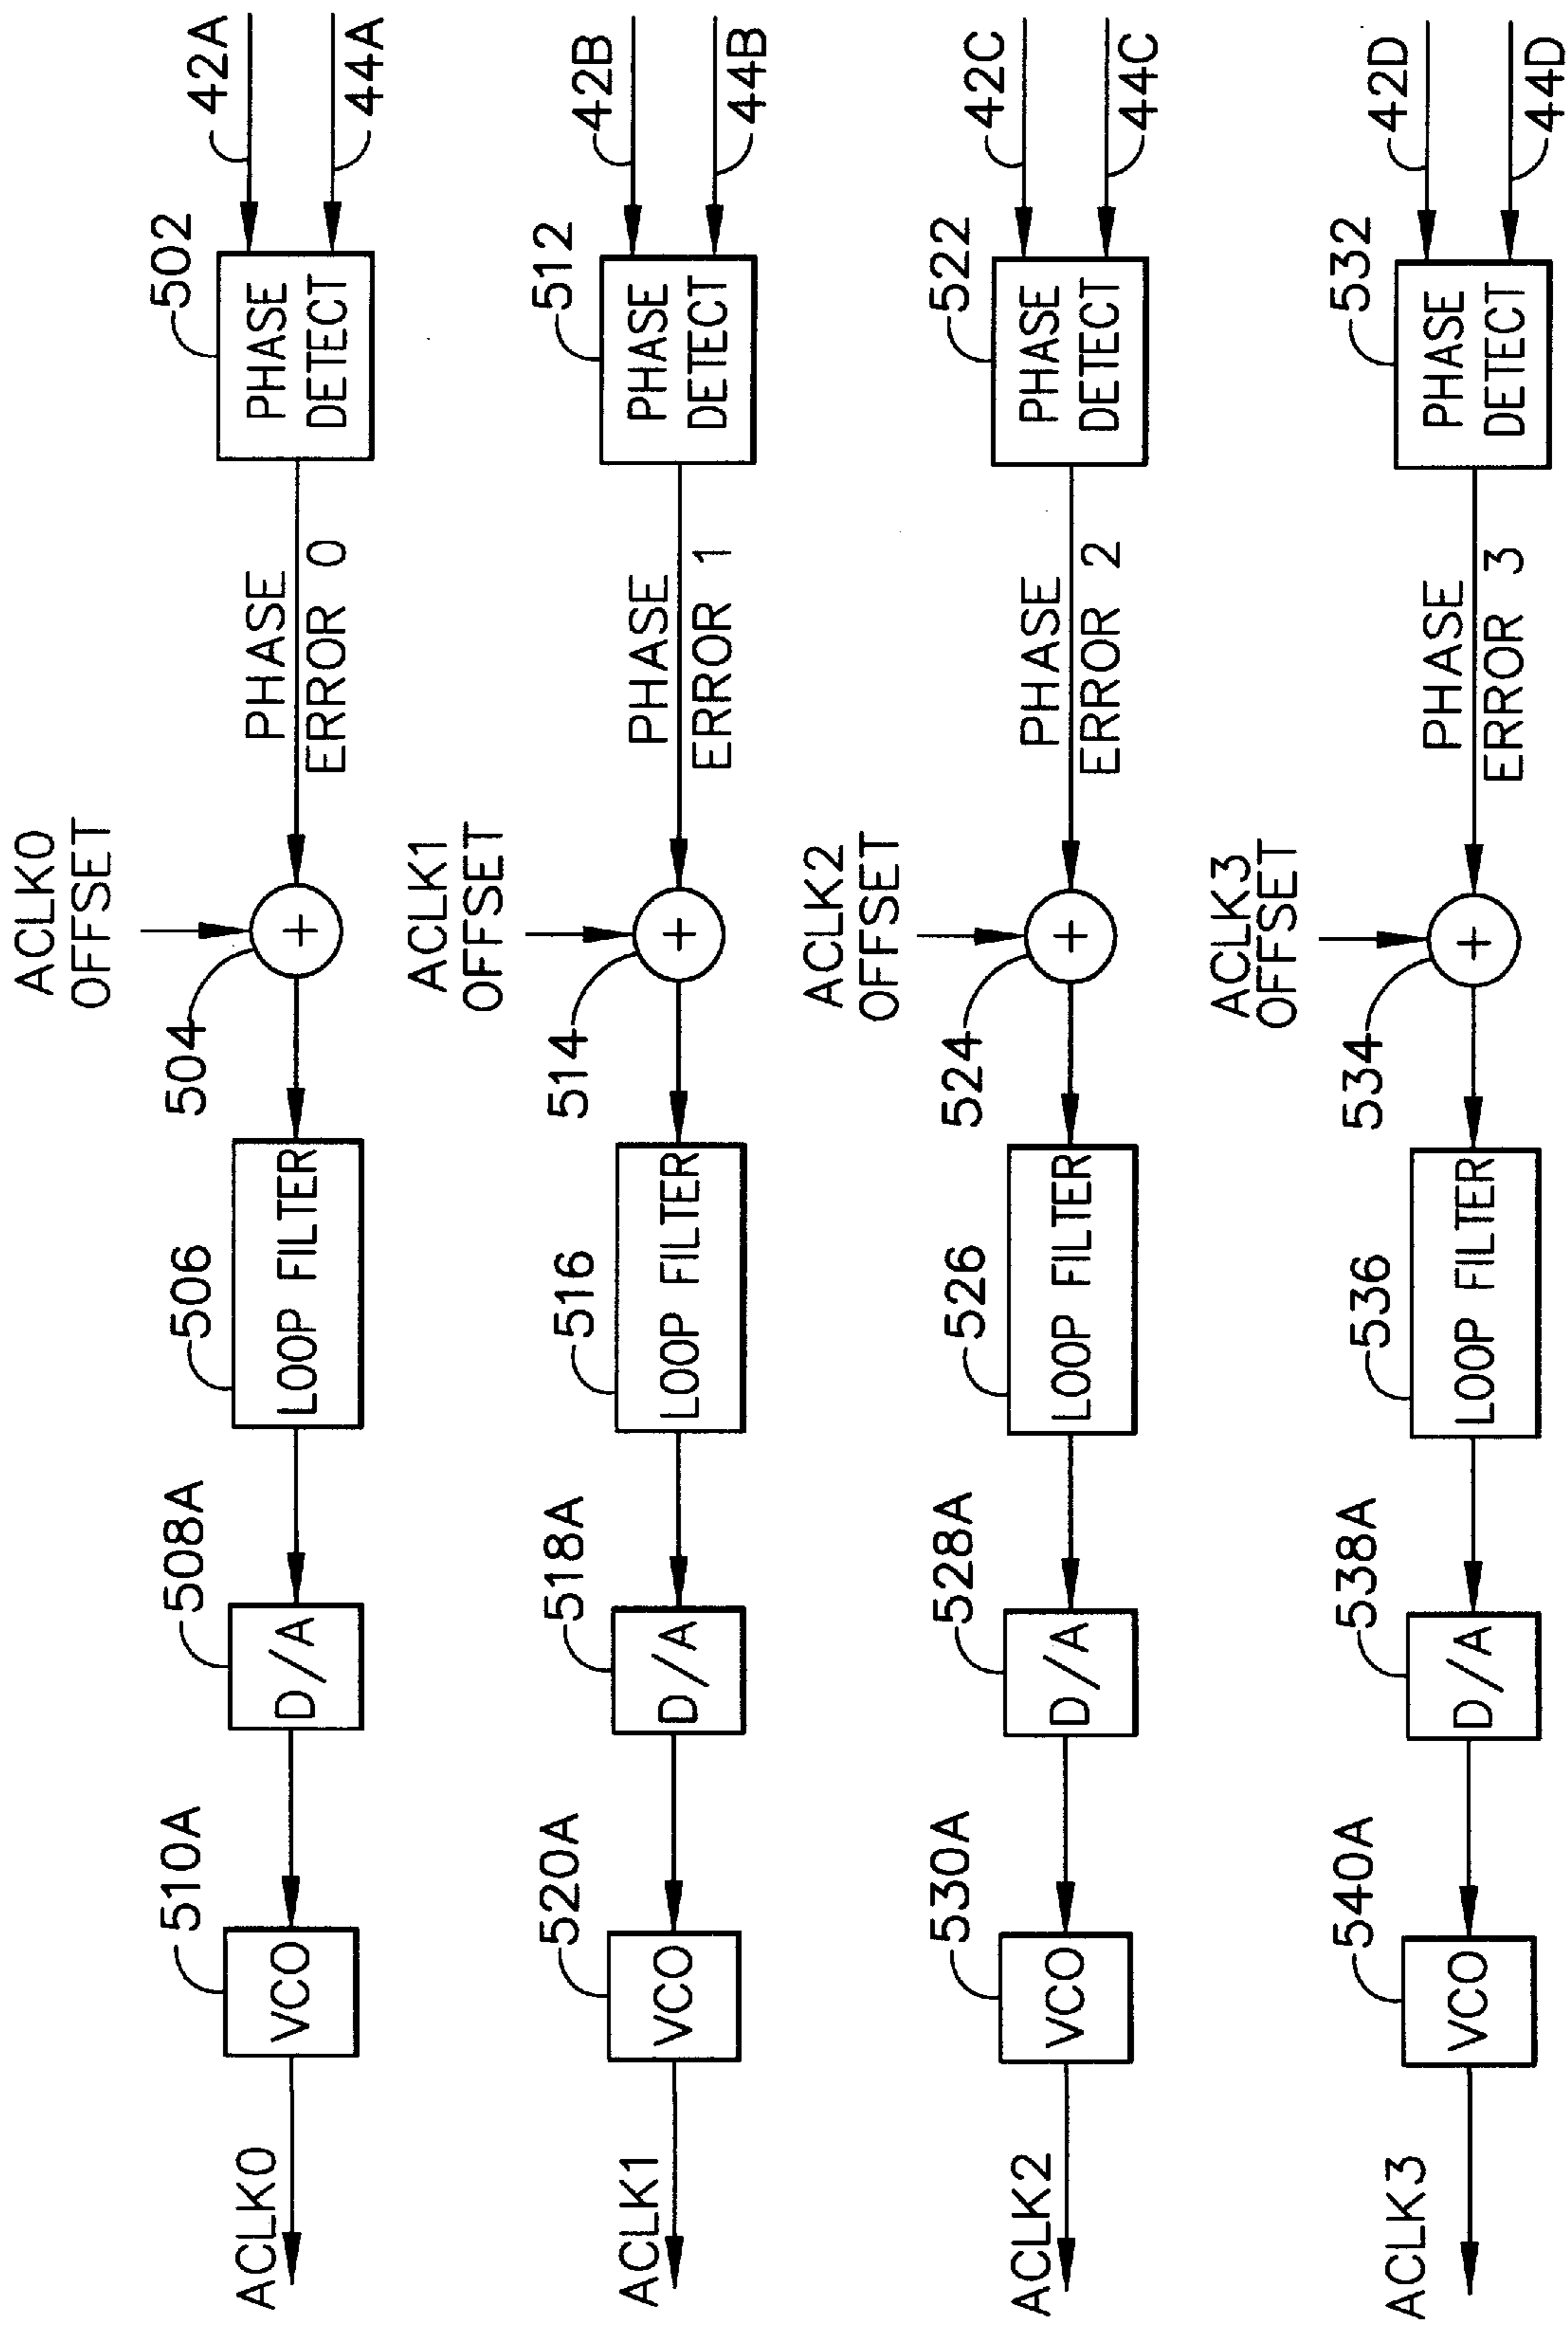


FIG. 5A

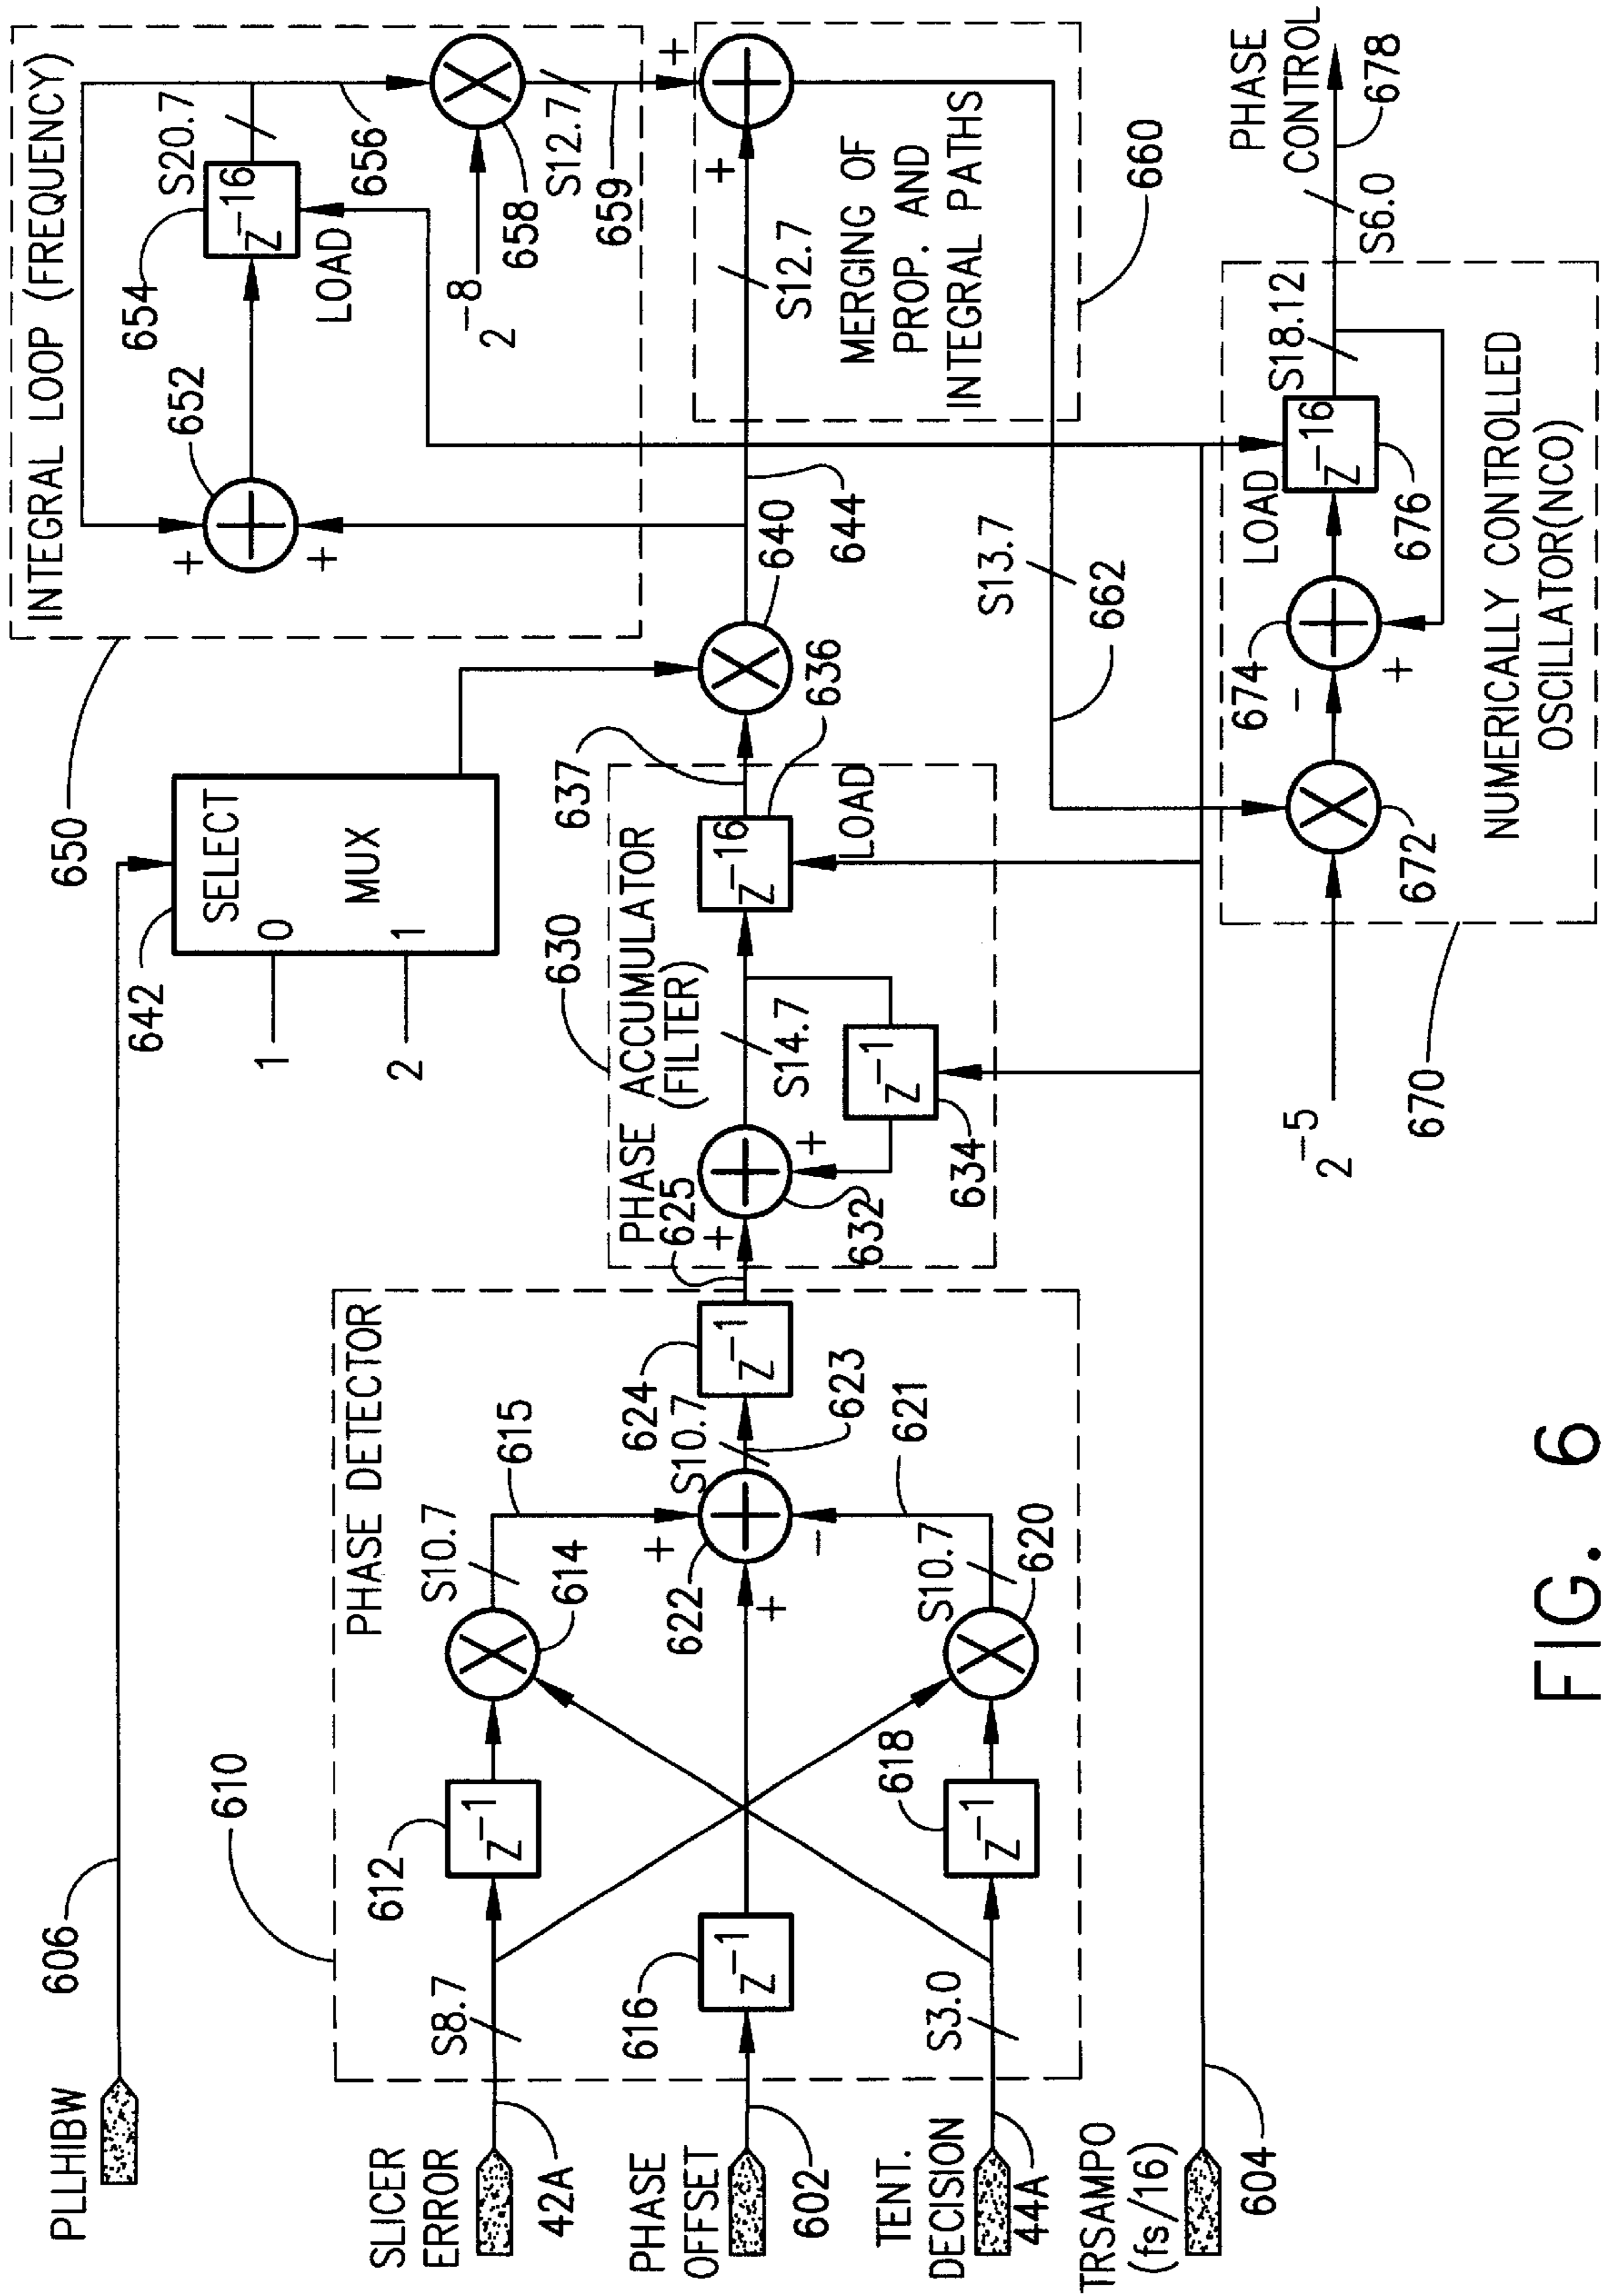


FIG. 6




FIG. 7




FIG. 8

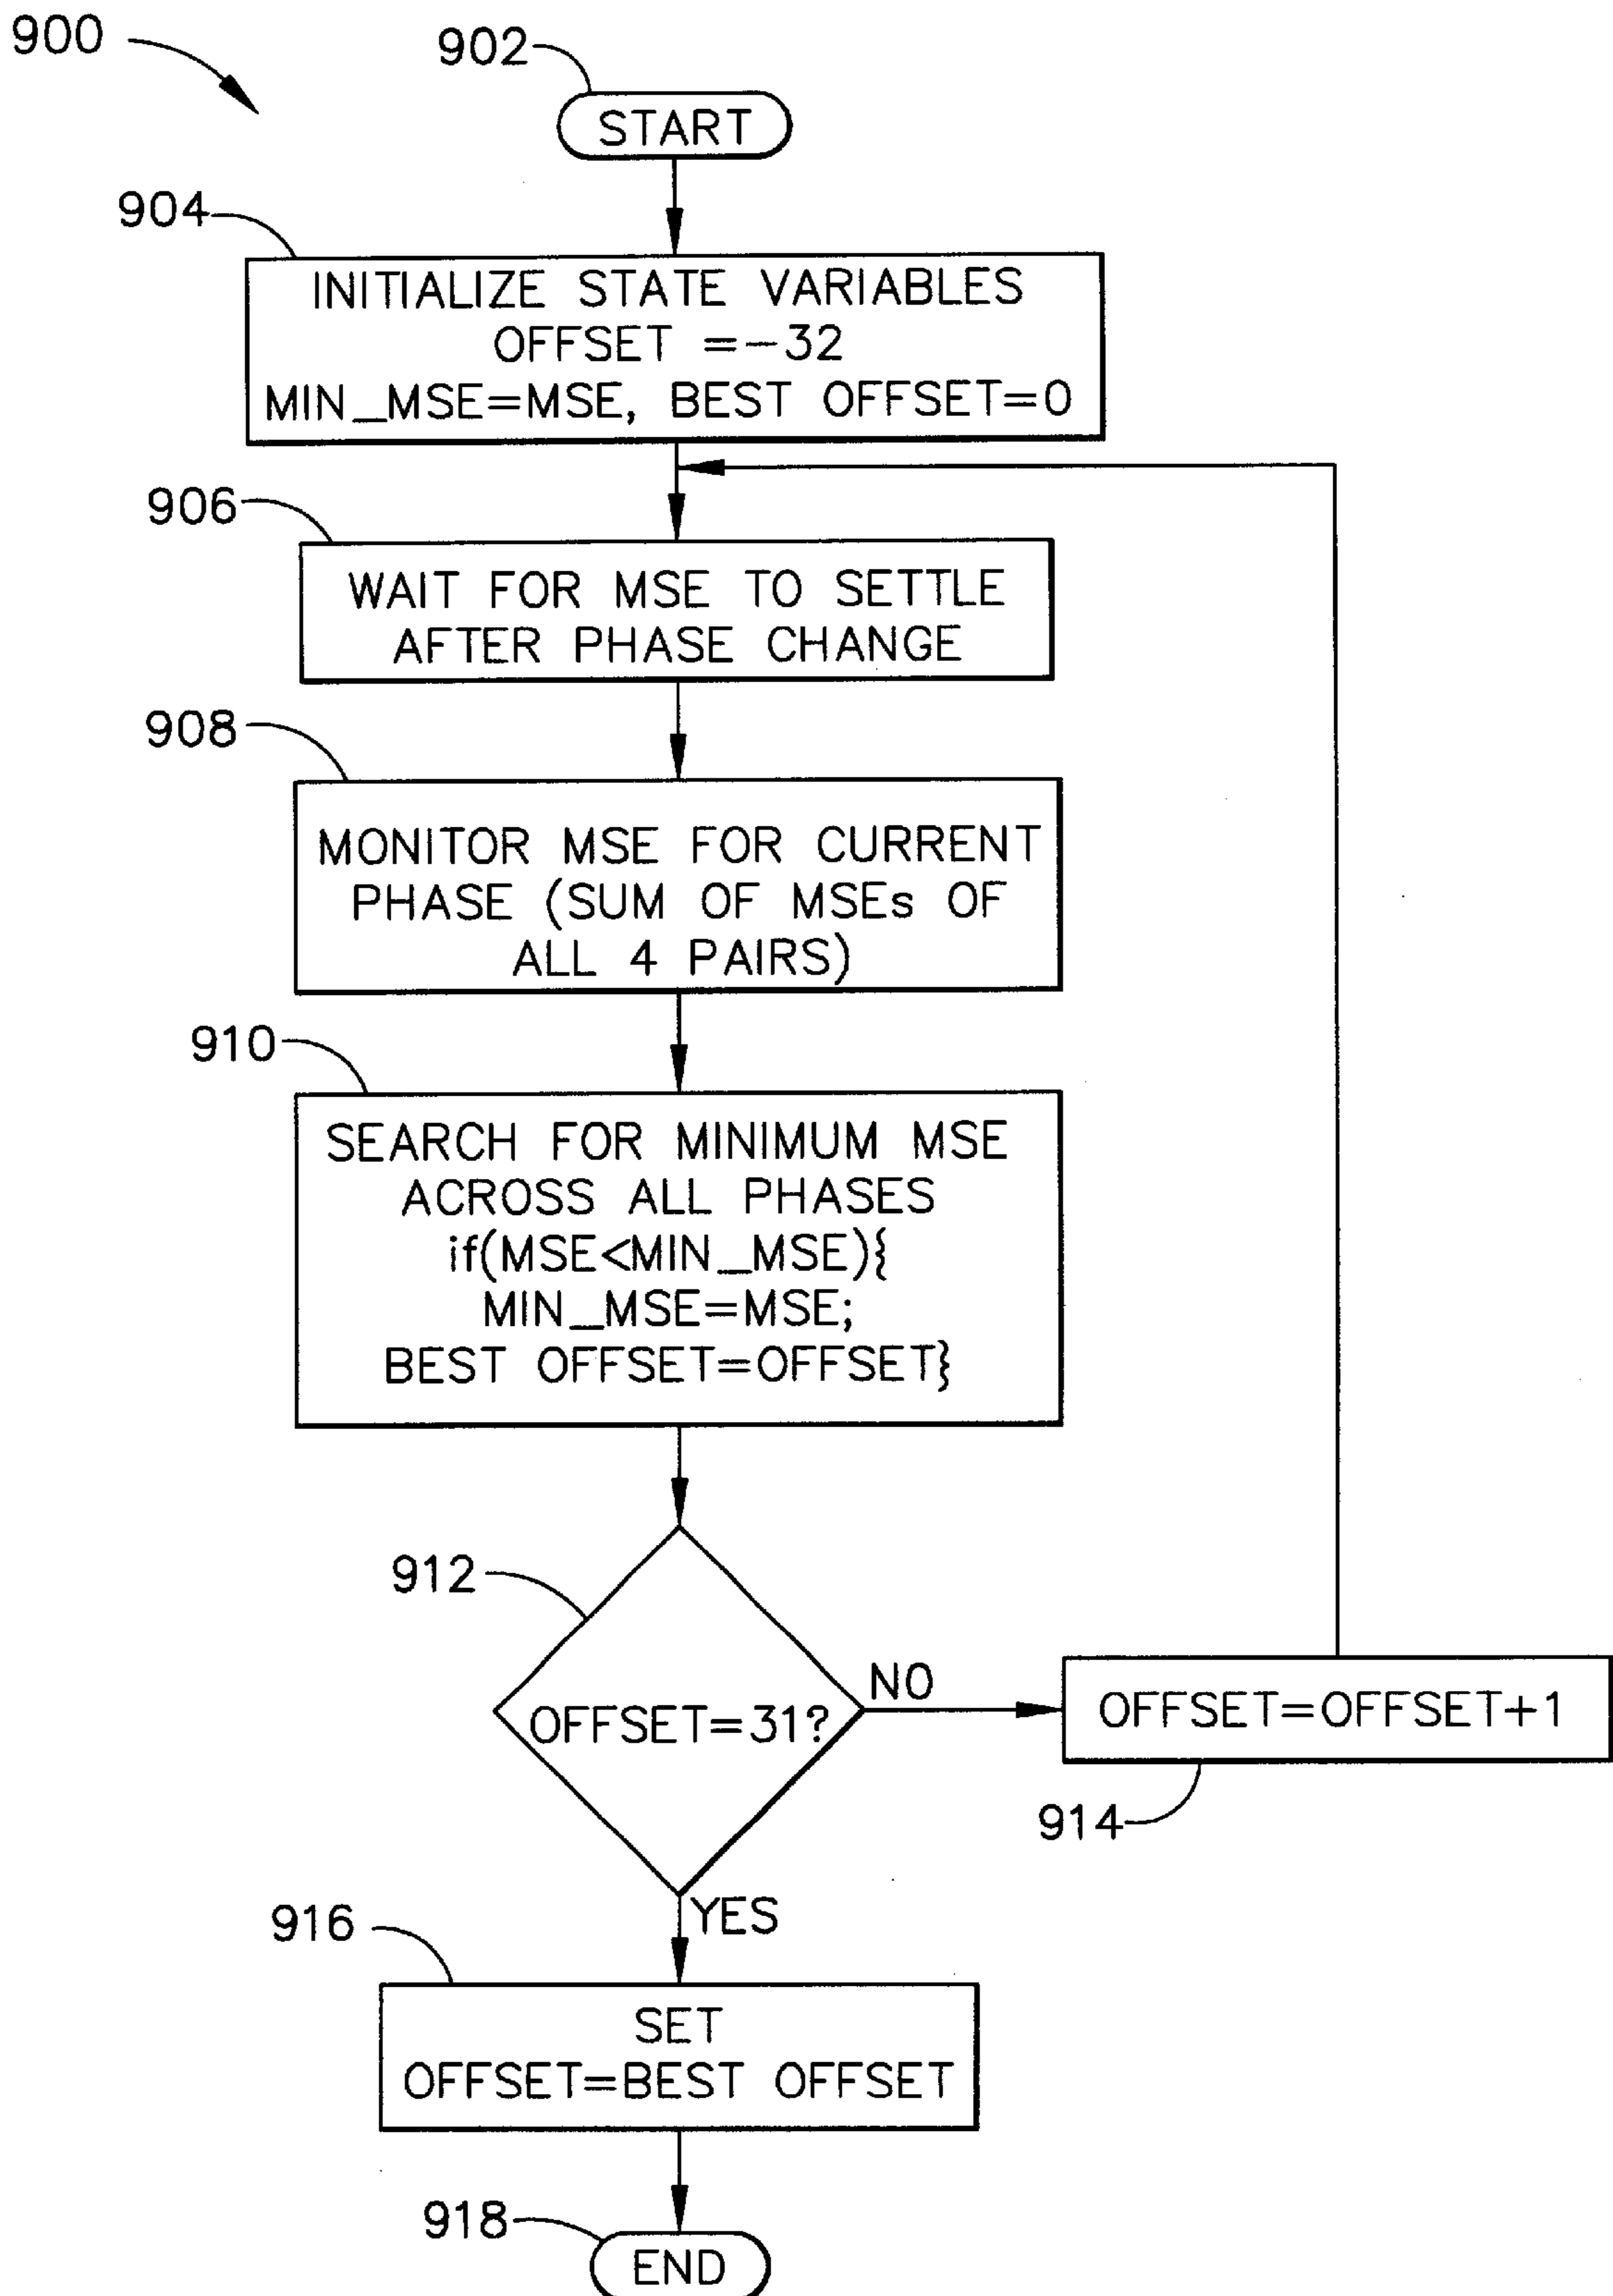


FIG. 9

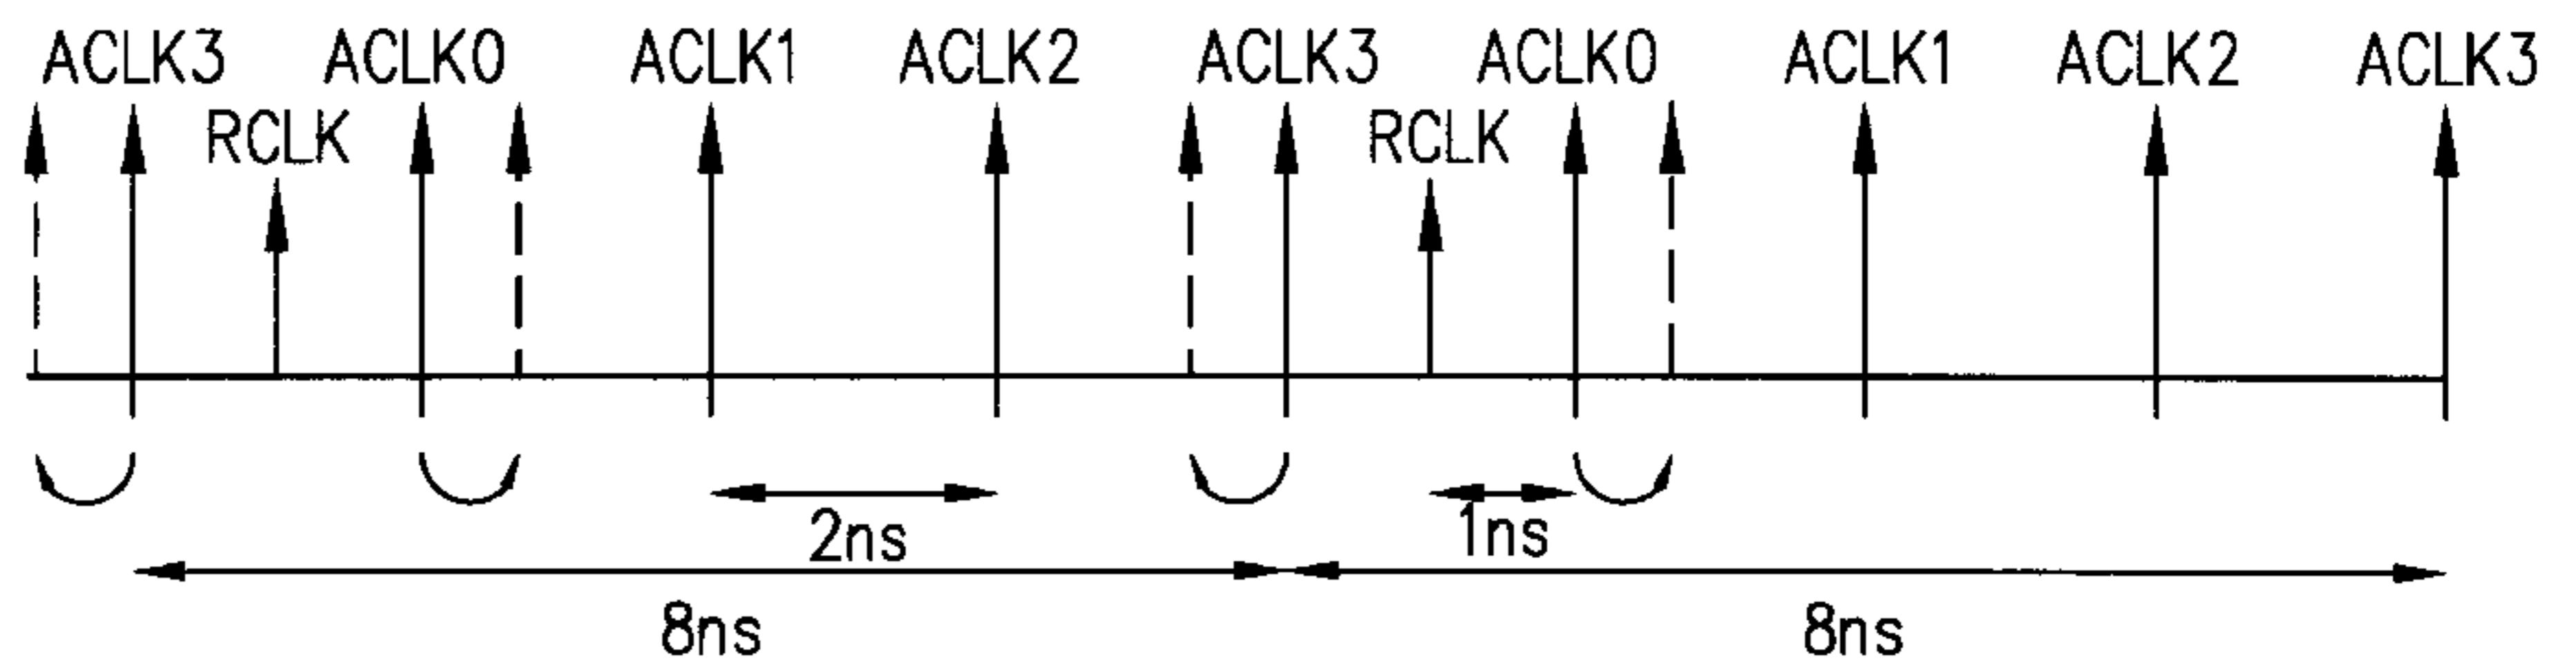


FIG. 10A




FIG. 10B




FIG. 10C

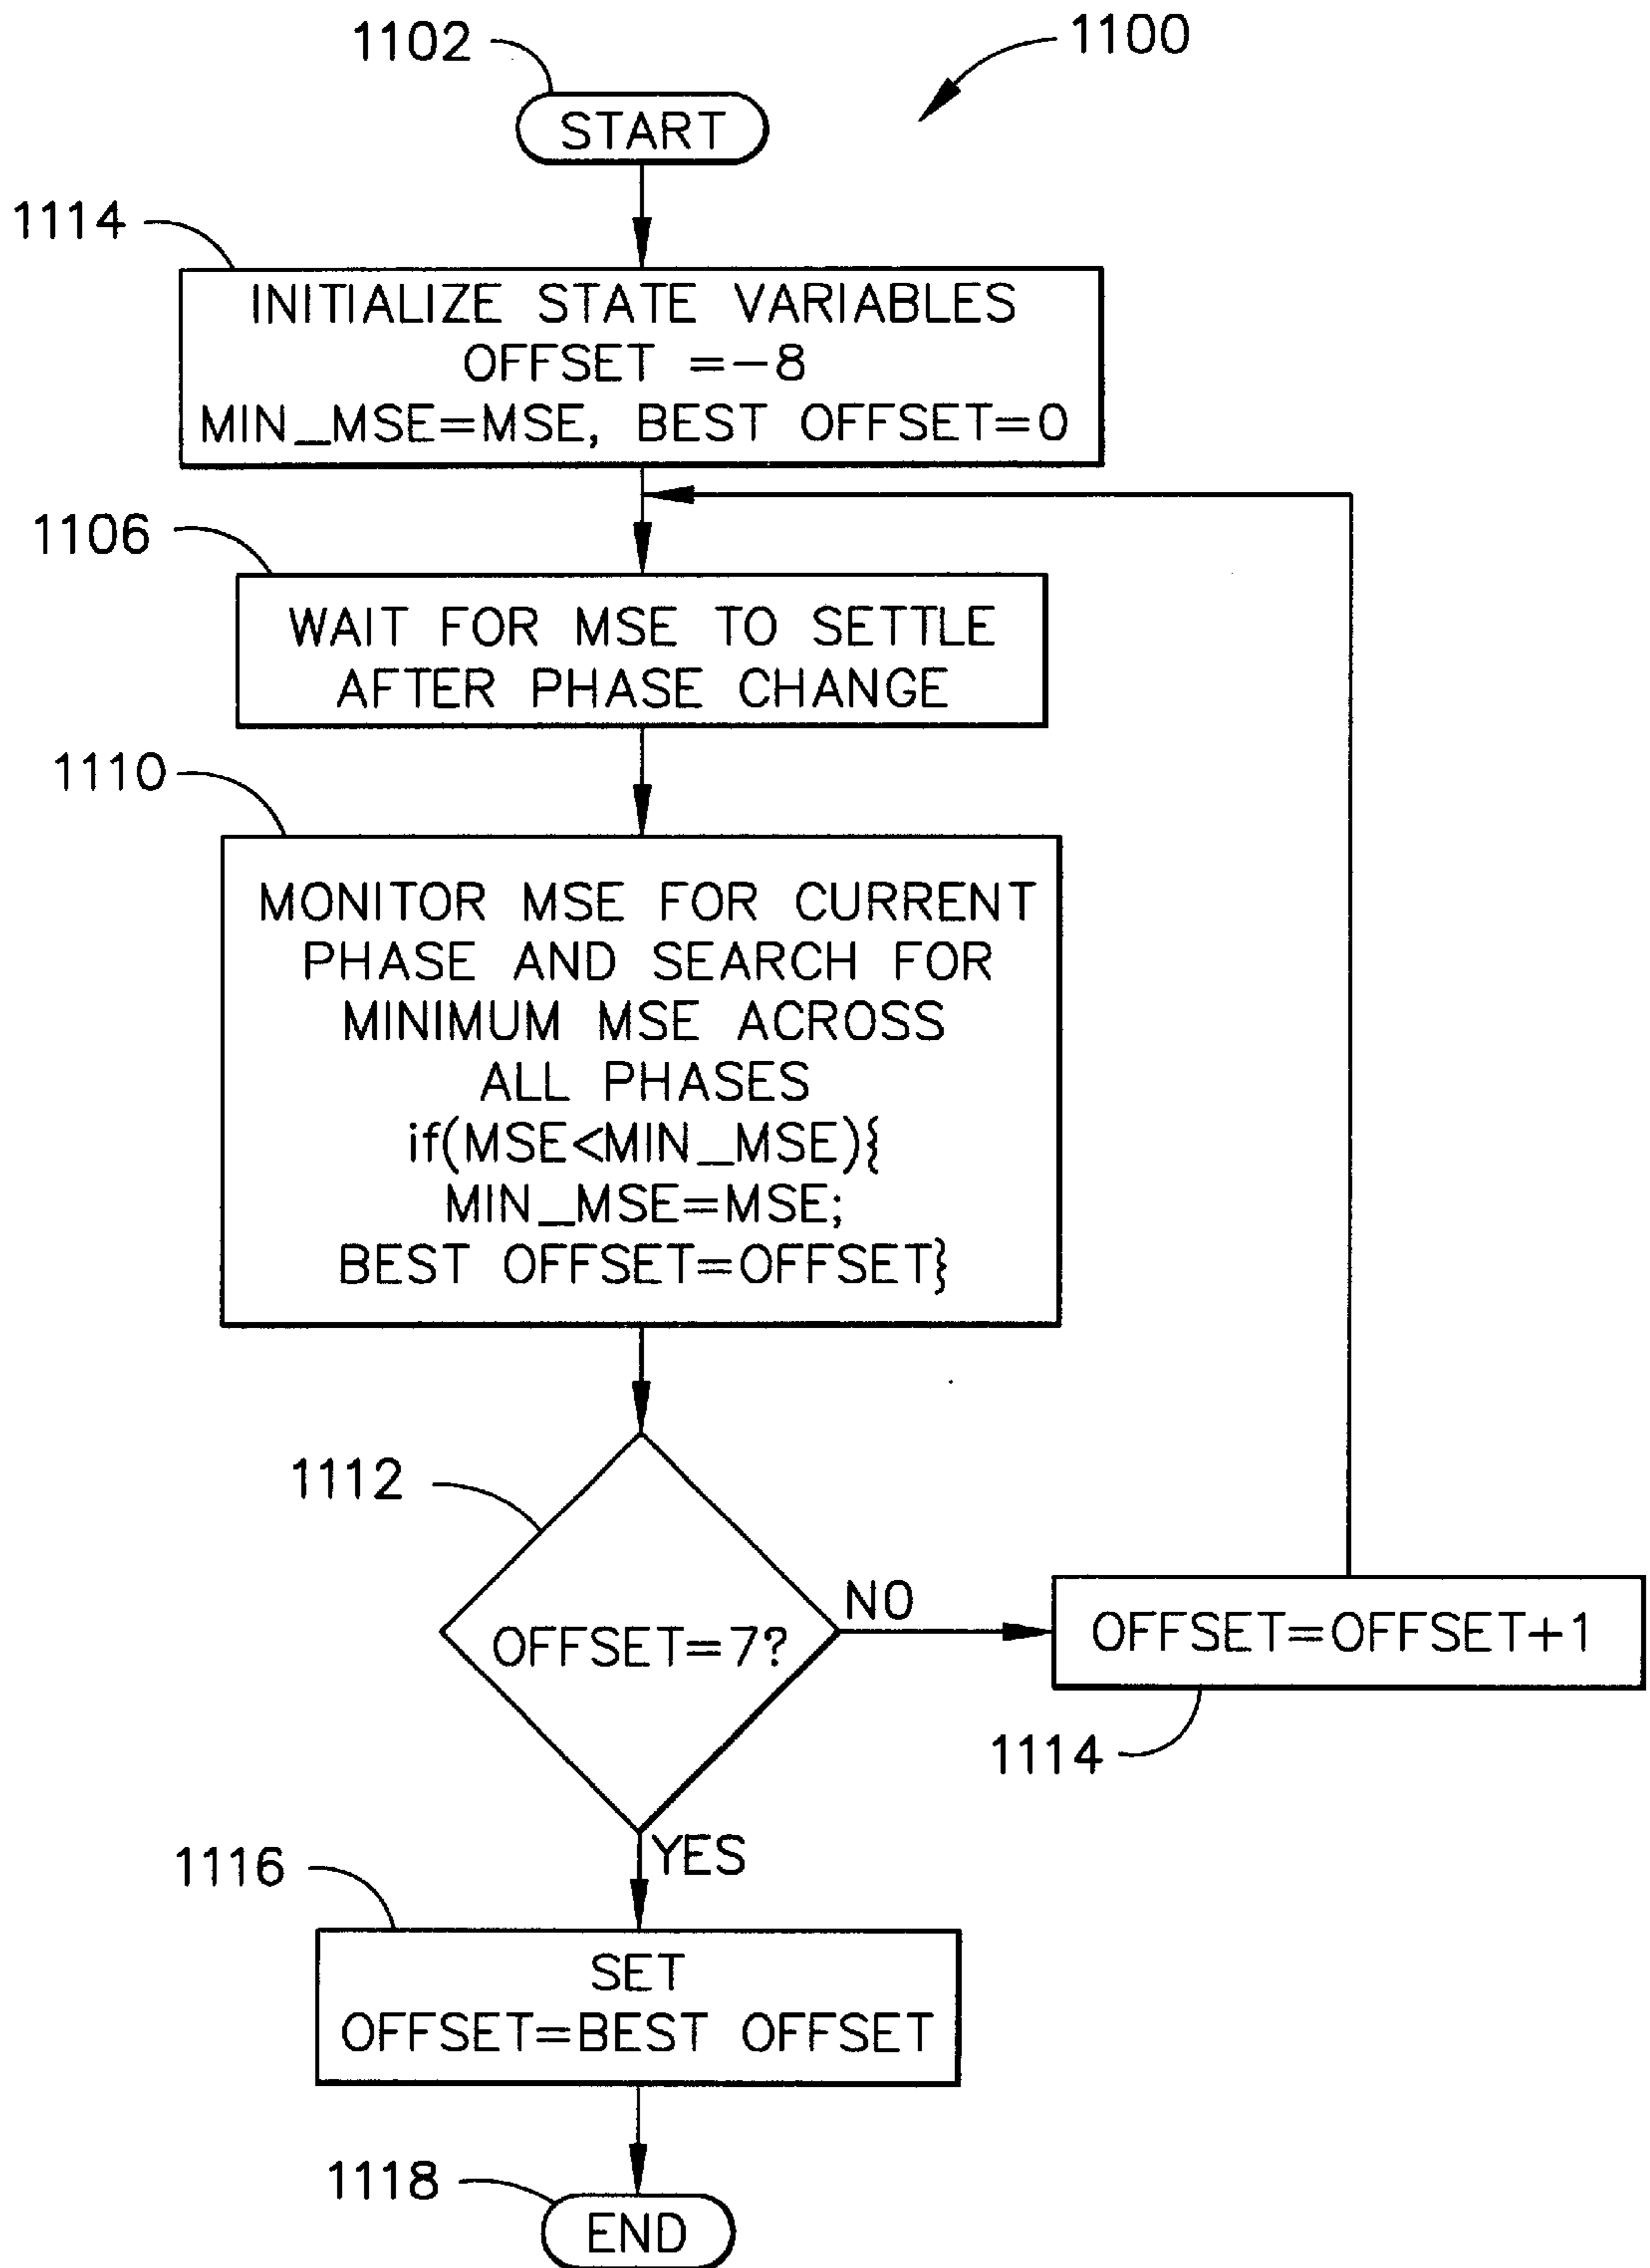


FIG. 11




FIG. 12

TIMING RECOVERY SYSTEM FOR A MULTI-PAIR GIGABIT TRANSCEIVER

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority of the following provisional applications, the contents of each of which are herein incorporated by reference: Ser. No. 60/107,874 entitled "Apparatus for, and Method of, Distributing Clock Signals in a Communications System" filed on Nov. 9, 1998; Ser. No. 60/108,319 entitled "Gigabit Ethernet Transceiver" filed on Nov. 13, 1998; Ser. No. 60/108,648 entitled "Clock Generation and Distribution in an Ethernet Transceiver" filed on Nov. 16, 1998 and Ser. No. 60/130,616 entitled "Multi-Pair Gigabit Ethernet Transceiver" filed on Apr. 22, 1999.

The present invention is related to the following co-pending applications filed on the same day as the present invention and assigned to the same assignee, the contents of each of which are herein incorporated by reference: Ser. No. 09/437,724 entitled "Switching Noise Reduction in a Multi-Clock Domain Transceiver" and Ser. No. 09/937,719 entitled "Multi-Pair Gigabit Ethernet Transceiver".

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention generally relates to clock signals in a transceiver. More particularly, the present invention relates to a method and an apparatus for generating and distributing clock signals in a gigabit Ethernet transceiver which includes more than one constituent transceiver.

2. Description of Related Art

A transceiver includes a transmitter and a receiver. In a traditional half-duplex transceiver, the transmitter and the receiver can operate with a common clock signal since the transmitting and receiving operations do not occur simultaneously.

In a full-duplex transceiver, the transmitting operation occurs simultaneously with the receiving operation. The full-duplex transceiver needs to operate with at least two clock signals, a transmit clock signal (TCLK) and a sampling clock signal. The TCLK signal is used by the transmitter to regulate transmission of data symbols. The sampling clock signal is used by the receiver to regulate sampling of the received signal at an analog-to digital (A/D) converter. At the local receiver, the frequency and phase of the sampling clock signal are adjusted by a timing recovery system of the local receiver in such a way that they track the transmit clock signal of the remote transmitter. The sampled received signal is demodulated by digital signal processing function blocks of the receiver. These digital processing functions blocks may operate in accordance with either the TCLK signal or the sampling clock signal, provided that signals crossing boundaries between the two clock signals are treated appropriately so that any loss of signal or data samples is prevented.

The IEEE 802.3ab standard (also called 1000BASE-T) for 1 gigabit per second (Gb/s) Ethernet full-duplex communication system specifies that there are four constituent transceivers in a gigabit transceiver and that the full-duplex communication is over four twisted pairs of Category-5 copper cables. Since a Gigabit Ethernet transceiver has four constituent transmitters and four constituent receivers, its operation is much more complex than the operation of a traditional full-duplex transceiver. The four twisted pairs of

cable may introduce different delays on the signals, causing the signals to have different phases. This, in turn, requires the gigabit Ethernet transceiver to have four A/D converters operating in accordance with four respective sampling clock signals. In addition, the problem of switching noise coupled from the digital signal processing blocks of the gigabit Ethernet transceiver to the four A/D converters must also be addressed.

Therefore, there is a need to have an efficient method and system for generating the clock signals for a gigabit Ethernet transceiver. There is also a need to distribute the clock signals such that effect of switching noise is minimized.

SUMMARY OF THE INVENTION

The present invention provides a method and a timing recovery system for generating a set of clock signals in a system which includes a set of subsystems. Each of the subsystems includes an analog section. The set of clock signals includes a set of sampling clock signals. Each of the analog sections operates in accordance with a corresponding one of the sampling clock signals. For each of the sampling clock signals, a phase error is generated from a corresponding phase detector. The phase errors are filtered by a set of corresponding loop filters. The filtered phase errors are provided to a set of corresponding oscillators to generate phase control signals. The phase control signals are provided to a set of corresponding phase selectors to generate the sampling clock signals.

BRIEF DESCRIPTION OF THE DRAWINGS

The features of the present invention will become more apparent and the invention will be best understood by reference to the following description and the accompanying drawings, wherein:

FIG. 1 a simplified block diagram of a multi-pair communication system operating in conformance with the IEEE 802.3ab standard (also termed 1000BASE-T) for 1 gigabit (Gb/s) Ethernet full-duplex communication over four twisted pairs of Category-5 copper wires;

FIG. 2 is a simplified block diagram of the functional architecture and internal construction of an embodiment of a gigabit transceiver of FIG. 1;

FIG. 3 is a simplified block diagram of an embodiment of the trellis decoder 38 of FIG. 2;

FIG. 4 illustrates the general clocking relationship between the transmitter and the receiver inside each of the four constituent transceivers 108 of the gigabit Ethernet transceiver (101 or 102) of FIG. 1;

FIG. 5 is a simplified block diagram of an embodiment of the timing recovery system constructed according to the present invention;

FIG. 6 is a block diagram of an exemplary implementation of the system of FIG. 5;

FIG. 7 is a block diagram of an exemplary embodiment of the phase reset logic block used for resetting the register of the NCO of FIG. 6 to a specified value;

FIG. 8 is a block diagram of an exemplary phase shifter logic block used for the phase control of the receive clock signal RCLK;

FIG. 9 is a flowchart of an embodiment of the process for adjusting the phase of the receive clock signal RCLK;

FIG. 10A is a first example of clock distribution where the transitions of the four sampling clock signals ACLK0-3 are evenly distributed within the symbol period.

FIG. 10B is a second example of clock distribution where the transitions of the four sampling clock signals ACLK0–3 are distributed within the symbol period of 8 nanoseconds (ns) such that each ACLK clock transition is 1 ns apart from an adjacent ACLK clock transition.

FIG. 10C is a third example of clock distribution where the transitions of the four sampling clock signals ACLK0–3 occur at the same instant within the symbol period.

FIG. 11 is a flowchart of an embodiment of the process for adjusting the phase of a sampling clock signal ACLKx associated with one of the constituent transceivers;

FIG. 12 is a block diagram of an embodiment of the MSE computation block used for computing the mean squared error of a constituent transceiver.

DETAILED DESCRIPTION OF THE INVENTION

The present invention provides a method and a timing recovery system for generating a set of clock signals in a processing system. The set of clock signals includes a set of sampling clock signals. The processing system includes a set of processing subsystems, each of which includes an analog section and a digital section. Each of the analog sections operates in accordance with a corresponding sampling clock signals. The digital sections operate in accordance with a receive clock. An example of the processing system is a gigabit transceiver. In this case, the processing subsystems are the constituent transceivers.

The present invention also provides a method and a system for substantially minimizing system performance degradation caused by coupling of switching noise from the digital sections to the analog sections.

The present invention can be used to generate and distribute clock signals in a gigabit transceiver of a Gigabit Ethernet communication system such that effect of switching noise coupled from one clock domain to another clock domain is minimized. By “clock domain”, it is meant the circuit blocks that are operating according to transitions of a particular clock signal. For ease of explanation, the present invention will be described in detail as applied to this exemplary application. However, this is not to be construed as a limitation of the present invention.

In order to appreciate the advantages of the present invention, it will be beneficial to describe the invention in the context of an exemplary bi-directional communication device, such as an Ethernet transceiver. The particular exemplary implementation chosen is depicted in FIG. 1, which is a simplified block diagram of a multi-pair communication system operating in conformance with the IEEE 802.3ab standard (also termed 1000BASE-T) for 1 gigabit (Gb/s) Ethernet full-duplex communication over four twisted pairs of Category-5 copper wires.

In FIG. 1, the communication system is represented as a point-to-point system in order to simplify the explanation, and includes two main transceiver blocks 102 and 104, coupled together via four twisted-pair cables 112a, b, c and d. Each of the wire pairs 112a, b, c, d is coupled to each of the transceiver blocks 102, 104 through a respective one of four line interface circuits 106. Each of the wire pairs 112a, b, c, d facilitates communication of information between corresponding pairs of four pairs of transmitter/receiver circuits (constituent transceivers) 108. Each of the constituent transceivers 108 is coupled between a respective line interface circuit 106 and a Physical Coding Sublayer (PCS) block 110. At each of the transceiver blocks 102 and 104, the four constituent transceivers 108 are capable of operating

simultaneously at 250 megabits of information data per second (Mb/s) each, i.e., 125 Mbaud at 2 information data bits per symbol, the 2 information data bits being encoded in one of the 5 levels of the PAM-5 (Pulse Amplitude Modulation) alphabet. The four constituent transceivers 108 are coupled to the corresponding remote constituent transceivers through respective line interface circuits to facilitate full-duplex bi-directional operation. Thus, 1Gb/s communication throughput of each of the transceiver blocks 102 and 104 is achieved by using four 250 Mb/s constituent transceivers 108 for each of the transceiver blocks 102, 104 and four pairs of twisted copper cables to connect the two transceiver blocks 102, 104 together.

FIG. 2 is a simplified block diagram of the functional architecture and internal construction of an exemplary transceiver block, indicated generally at 200, such as transceiver 101 of FIG. 1. Since the illustrative transceiver application relates to gigabit Ethernet transmission, the transceiver will be referred to as the “gigabit transceiver”. For ease of illustration and description, FIG. 2 shows only one of the four 250 Mb/s constituent transceivers which are operating simultaneously (termed herein 4-D operation). However, since the operation of the four constituent transceivers are necessarily interrelated, certain blocks and signal lines in the exemplary embodiment of FIG. 2 perform four-dimensional operations and carry four-dimensional (4-D) signals, respectively. By 4-D, it is meant that the data from the four constituent transceivers are used simultaneously. In order to clarify signal relationships in FIG. 2, thin lines correspond to 1-dimensional functions or signals (i.e., relating to only a single constituent transceiver), and thick lines correspond to 4-D functions or signals (relating to all four constituent transceivers).

Referring to FIG. 2, the gigabit transceiver 200 includes a Gigabit Medium Independent Interface (GMII) block 202 subdivided into a receive GMII circuit 202R and a transmit GMII circuit 202T. The transceiver also includes a Physical Coding Sublayer (PCS) block 204, subdivided into a receive PCS circuit 204R and a transmit PCS circuit 204T, a pulse shaping filter 206, a digital-to analog (D/A) converter block 208, and a line interface block 210, all generally encompassing the transmitter portion of the transceiver.

The receiver portion generally includes a highpass filter 212, a programmable gain amplifier (PGA) 214, an analog-to-digital (A/D) converter 216, an automatic gain control (AGC) block 220, a timing recovery block 222, a pair-swap multiplexer block 224, a demodulator 226, an offset canceller 228, a near-end crosstalk (NEXT) canceller block 230 having three constituent NEXT cancellers and an echo canceller 232.

The gigabit transceiver 200 also includes an A/D first-in-first-out buffer (FIFO) 218 to facilitate proper transfer of data from the analog clock region to the receive clock region, and a loopback FIFO block (LPBK) 234 to facilitate proper transfer of data from the transmit clock region to the receive clock region. The gigabit transceiver 200 can optionally include an additional adaptive filter to cancel far-end crosstalk noise (FEXT canceller).

In operational terms, on the transmit path, the transmit section 202T of the GMII block receives data from the Media Access Control (MAC) module in byte-wide format at the rate of 125 MHz and passes them to the transmit section 204T of the PCS block via the FIFO 201. The FIFO 201 ensures proper data transfer from the MAC layer to the Physical Coding (PHY) layer, since the transmit clock of the PHY layer is not necessarily synchronized with the clock of

the MAC layer. In one embodiment, this small FIFO **201** has from about three to about five memory cells to accommodate the elasticity requirement which is a function of frame size and frequency offset.

The PCS transmit section **204T** performs certain scrambling operations and, in particular, is responsible for encoding digital data into the requisite codeword representations appropriate for transmission. In the illustrated embodiment of FIG. 2, the transmit PCS section **204T** incorporates a coding engine and signal mapper that implements a trellis coding architecture, such as required by the IEEE 802.3ab specification for gigabit transmission.

In accordance with this encoding architecture, the PCS transmit section **204T** generates four 1-D symbols, one for each of the four constituent transceivers. The 1-D symbol generated for the constituent transceiver depicted in FIG. 2 is filtered by the pulse shaping filter **206**. This filtering assists in reducing the radiated emission of the output of the transceiver such that it falls within the parameters required by the Federal Communications Commission. The pulse shaping filter **206** is implemented so as to define a transfer function of $0.75+0.25z^{-1}$. This particular implementation is chosen so that the power spectrum of the output of the transceiver falls below the power spectrum of a 100Base-Tx signal. The 100Base-Tx is a widely used and accepted Fast Ethernet standard for 100 Mb/s operation on two pairs of Category-5 twisted pair cables. The output of the pulse shaping filter **206** is converted to an analog signal by the D/A converter **208** operating at 125 MHz. The analog signal passes through the line interface block **210**, and is placed on the corresponding twisted pair cable.

On the receive path, the line interface block **210** receives an analog signal from the twisted pair cable. The received analog signal is preconditioned by the highpass filter **212** and the PGA **214** before being converted to a digital signal by the A/D converter **216** operating at a sampling rate of 125 MHz. The timing of the A/D converter **216** is controlled by the output of the timing recovery block **222**. The resulting digital signal is properly transferred from the analog clock region to the receive clock region by the A/D FIFO **218**. The output of the A/D FIFO **218** is also used by the AGC **220** to control the operation of the PGA **214**.

The output of the A/D FIFO **218**, along with the outputs from the A/D FIFOs of the other three constituent transceivers are inputted to the pair-swap multiplexer block **224**. The pair-swap multiplexer block **224** uses the 4-D pair-swap control signal from the receive section **204R** of PCS block to sort out the four input signals and send the correct signals to the respective feedforward equalizers **26** of the demodulator **226**. This pair-swapping control is needed for the following reason. The trellis coding methodology used for the gigabit transceivers (**101** and **102** of FIG. 1) is based on the fact that a signal on each twisted pair of wire corresponds to a respective 1-D constellation, and that the signals transmitted over four twisted pairs collectively form a 4-D constellation. Thus, for the decoding to work, each of the four twisted pairs must be uniquely identified with one of the four dimensions. Any undetected swapping of the four pairs would result in erroneous decoding. In an alternate embodiment of the gigabit transceiver, the pair-swapping control is performed by the demodulator **226**, instead of the combination of the PCS receive section **204R** and the pair-swap multiplexer block **224**.

The demodulator **226** includes a feed-forward equalizer (FFE) **26** for each constituent transceiver, coupled to a deskew memory circuit **36** and a decoder circuit **38**, imple-

mented in the illustrated embodiment as a trellis decoder. The deskew memory circuit **36** and the trellis decoder **38** are common to all four constituent transceivers. The FFE **26** receives the received signal intended for it from the pair-swap multiplexer block **224**. The FFE **26** is suitably implemented to include a precursor filter **28**, a programmable inverse partial response (IPR) filter **30**, a summing device **32**, and an adaptive gain stage **34**. The FFE **26** is a least-mean-squares (LMS) type adaptive filter which is configured to perform channel equalization as will be described in greater detail below.

The precursor filter **28** generates a precursor to the input signal **2**. This precursor is used for timing recovery. The transfer function of the precursor filter **28** might be represented as $-\gamma+z^{-1}$, with γ equal to $\frac{1}{16}$ for short cables (less than 80 meters) and $\frac{1}{8}$ for long cables (more than 80 m). The determination of the length of a cable is based on the gain of the coarse PGA **14** of the programmable gain block **214**.

The programmable IPR filter **30** compensates the ISI (intersymbol interference) introduced by the partial response pulse shaping in the transmitter section of a remote transceiver which transmitted the analog equivalent of the digital signal **2**. The transfer function of the IPR filter **30** may be expressed as $1/(1+Kz^{-1})$. In the present example, K has an exemplary value of 0.484375 during startup, and is slowly ramped down to zero after convergence of the decision feedback equalizer included inside the trellis decoder **38**. The value of K may also be any positive value strictly less than 1.

The summing device **32** receives the output of the IPR filter **30** and subtracts therefrom adaptively derived cancellation signals received from the adaptive filter block, namely signals developed by the offset canceller **228**, the NEXT cancellers **230**, and the echo canceller **232**. The offset canceller **228** is an adaptive filter which generates an estimate of signal offset introduced by component circuitry of the transceiver's analog front end, particularly offsets introduced by the PGA **214** and the A/D converter **216**.

The three NEXT cancellers **230** may also be described as adaptive filters and are used, in the illustrated embodiment, for modeling the NEXT impairments in the received signal caused by interference generated by symbols sent by the three local transmitters of the other three constituent transceivers. These impairments are recognized as being caused by a crosstalk mechanism between neighboring pairs of cables, thus the term near-end crosstalk, or NEXT. Since each receiver has access to the data transmitted by the other three local transmitters, it is possible to approximately replicate the NEXT impairments through filtering. Referring to FIG. 2, the three NEXT cancellers **230** filter the signals sent by the PCS block to the other three local transmitters and produce three signals replicating the respective NEXT impairments. By subtracting these three signals from the output of the IPR filter **30**, the NEXT impairments are approximately cancelled.

Due to the bi-directional nature of the channel, each local transmitter causes an echo impairment on the received signal of the local receiver with which it is paired to form a constituent transceiver. In order to remove this impairment, an echo canceller **232** is provided, which may also be characterized as an adaptive filter, and is used, in the illustrated embodiment, for modeling the signal impairment due to echo. The echo canceller **232** filters the signal sent by the PCS block to the local transmitter associated with the receiver, and produces an approximate replica of the echo impairment. By subtracting this replica signal from the

output of the IPR filter **30**, the echo impairment is approximately cancelled.

The adaptive gain stage **34** receives the processed signal from the summing circuit **32** and fine tunes the signal path gain using a zero-forcing LMS algorithm. Since this adaptive gain stage **34** trains on the basis of error signals generated by the adaptive filters **228**, **230** and **232**, it provides a more accurate signal gain than the one provided by the PGA **214** in the analog section.

The output of the adaptive gain stage **34**, which is also the output of the FFE **26**, is inputted to the deskew memory circuit **36**. The deskew memory **36** is a four-dimensional function block, i.e., it also receives the outputs of the three FFEs of the other three constituent transceivers. There may be a relative skew in the outputs of the four FFEs, which are the four signal samples representing the four symbols to be decoded. This relative skew can be up to 50 nanoseconds, and is due to the variations in the way the copper wire pairs are twisted. In order to correctly decode the four symbols, the four signal samples must be properly aligned. The deskew memory aligns the four signal samples received from the four FFEs, then passes the deskewed four signal samples to a decoder circuit **38** for decoding.

In the context of the exemplary embodiment, the data received at the local transceiver was encoded before transmission, at the remote transceiver. In the present case, data might be encoded using an 8-state four-dimensional trellis code, and the decoder **38** might therefore be implemented as a trellis decoder. In the absence of intersymbol interference (ISI), a proper 8-state Viterbi decoder would provide optimal decoding of this code. However, in the case of Gigabit Ethernet, the Category-5 twisted pair cable introduces a significant amount of ISI. In addition, the partial response filter of the remote transmitter on the other end of the communication channel also contributes some ISI. Therefore, the trellis decoder **38** must decode both the trellis code and the ISI, at the high rate of 125 MHz. In the illustrated embodiment of the gigabit transceiver, the trellis decoder **38** includes an 8-state Viterbi decoder, and uses a decision-feedback sequence estimation approach to deal with the ISI components.

The 4-D output of the trellis decoder **38** is provided to the PCS receive section **204R**. The receive section **204R** of the PCS block de-scrambles and decodes the symbol stream, then passes the decoded packets and idle stream to the receive section **202R** of the GMII block which passes them to the MAC module. The 4-D outputs, which are the error and tentative decision, respectively, are provided to the timing recovery block **222**, whose output controls the sampling time of the A/D converter **216**. One of the four components of the error and one of the four components of the tentative decision correspond to the receiver shown in FIG. 2, and are provided to the adaptive gain stage **34** of the FFE **26** to adjust the gain of the equalizer signal path. The error component portion of the decoder output signal is also provided, as a control signal, to adaptation circuitry incorporated in each of the adaptive filters **230** and **232**. Adaptation circuitry is used for the updating and training process of filter coefficients.

FIG. 3 is a block diagram of the trellis decoder **38** of FIG. 2. The trellis decoder **38** includes a multiple decision feedback equalizer (MDFE) **302**, a Viterbi decoder **304**, a path metrics module **306**, a path memory module **308**, a select logic **310**, and a decision feedback equalizer **312**.

The Viterbi decoder **304** performs 4D slicing of the Viterbi inputs provided by the MDFE **302** and computes the

branch metrics. Based on the branch metrics and the previous path metrics received from the path metrics module **306**, the Viterbi decoder **304** extends the paths and computes the extended path metrics. The Viterbi decoder **304** selects the best path incoming to each of the 8 states, updates the path memory stored in the path memory module **308** and the path metrics stored in the path metrics module **306**.

The computation of the final decision and the tentative decisions are performed in the path memory module **308** based on the 4D symbols stored in the path memory for each state. At each iteration of the Viterbi algorithm, the best of the 8 states, i.e., the one associated with the path having the lowest path metric, is selected, and the 4D symbol from the associated path stored at the last level of the path memory is selected as the final decision **40** and provided to the receive section of the PCS **204R** (FIG. 2). Symbols at lower depth levels are selected as tentative decisions, which are used to feed the delay line of the DFE **312**.

The number of the outputs V_i to be used as tentative decisions depends on the required accuracy and speed of decoding operation. A delayed version of V_{OF} is provided as the 4D tentative decision **44** (FIG. 2) to the Feed-Forward Equalizers **26** of the 4 constituent transceivers and the timing recovery block **222** (FIG. 2).

Based on the symbols V_{OF} , V_{1F} , and V_{2F} , the DFE **612** produces the intersymbol interference (ISI) replica associated with all previous symbols except the two most recent (since it was derived without using the first two taps of the DFE **612**). The ISI replica is fed to the MDFE **302** (this ISI replica is denoted as the "tail component" in FIG. 6). The MDFE **302** computes the ISI replica associated with all previous symbols including the two most recent symbols, subtracts it from the output **37** of the deskew memory block **36** (FIG. 2) and provides the resulting Viterbi inputs to the Viterbi decoder **304**.

The DFE **312** also computes an ISI replica associated with the two most recent symbols, based on tentative decisions V_{OF} , V_{1F} , and V_{2F} . This ISI replica is subtracted from a delayed version of the output **37** of the de-skew memory block **36** to provide the soft decision **43**. The tentative decision V_{OF} is subtracted from the soft decision **43** to provide the error **42**. There are 3 different versions of the error **42**, which are 42_{enc} , 42_{ph} and 42_{dfe} . The error 42_{enc} is provided to the echo cancellers and NEXT cancellers of the constituent transceivers. The error 42_{ph} is provided to the FFEs **26** (FIG. 2) of the 4 constituent transceivers and the timing recovery block **222**. The error 42_{dfe} is used for the adaptation of the coefficients of the DFE **312**. The tentative decision **44** shown in FIG. 3 is a delayed version of V_{OF} . The soft decision **43** is only used for display purposes.

For the exemplary gigabit transceiver system **200** described above and shown in FIGS. 2 and 3, there is a PHY Control system (not shown) which provides control signals to the blocks shown in FIG. 2, including the timing recovery block **222**, to control their functions.

For the exemplary gigabit transceiver system **200** described above and shown in FIG. 2, there are design considerations regarding the allocation of boundaries of the clock domains. These design considerations are dependent on the clocking relationship between transmitters and receivers in a gigabit transceiver. Therefore, this clocking relationship will be discussed first.

During a bidirectional communication between two gigabit transceivers **101**, **102** (FIG. 1), through a process called "auto-negotiation", one of the gigabit transceivers assumes the role of the master while the other assumes the role of the

slave. When a gigabit transceiver assumes one of the two roles with respect to the remote gigabit transceiver, each of its constituent transceivers assumes the same role with respect to the corresponding one of the remote constituent transceivers. Each constituent transceiver **108** is constructed such that it can be dynamically configured to act as either the master or the slave with respect to a remote constituent transceiver **108** during a bi-directional communication. The clocking relationship between the transmitter and receiver inside the constituent transceiver **108** depends on the role of the constituent transceiver (i.e., master or slave) and is different for each of the two cases.

FIG. 4 illustrates the general clocking relationship on the conceptual level between the transmitter and the receiver of the gigabit Ethernet transceiver (**101** or **102**) of FIG. 1. For this conceptual FIG. 4, the transmitter TX represents the four constituent transmitters and the receiver RX represents the four constituent receivers.

Referring to FIG. 4, the gigabit transceiver **401** acts as the master while the gigabit transceiver **402** acts as the slave. The master **401** includes a transmitter **410** and a receiver **412**. The slave **402** includes a transmitter **420** and a receiver **422**. The transceiver **401** (respectively, **402**) receives from the GMII **202T** (FIG. 2) the data to be transmitted TXD via its input **413** (respectively, **423**), and the GMII transmit clock GTX_CLK (this clock is also called “gigabit transmit clock” in the IEEE 802.3ab standard) via its input **415** (respectively, **425**). The transceiver **401** (respectively, **402**) sends to the GMII **202R** (FIG. 2) the received data RXD via its output **417** (respectively, **427**), and the GMII receive clock RX_CLK (this clock is also called “gigabit receive clock” in the IEEE 802.3ab standard) via its output **419** (respectively, **429**). It is noted that the clocks GTX_CLK and RX_CLK may be different from the transmit clock TCLK and receive clock RCLK, respectively, of a gigabit transceiver.

The receiver **422** of the slave **402** synchronizes its receive clock to the transmit clock of the transmitter **410** of the master **401** in order to properly receive the data transmitted by the transmitter **410**. The transmit clock of the transmitter **420** of the slave **402** is essentially the same as the receive clock of the receiver **422**, thus it is also synchronized to the transmit clock of the transmitter **410** of the master **401**.

The receiver **412** of the master **401** is synchronized to the transmit clock of the transmitter **420** of the slave **402** in order to properly receive data sent by the transmitter **420**. Because of the synchronization of the receive and transmit clocks of the slave **402** to the transmit clock of transmitter **410** of the master **401**, the receive clock of the receiver **412** is synchronized to the transmit clock of the transmitter **410** with a phase delay (due to the twisted pairs of cables). Thus, in the absence of jitter, after synchronization, the receive clock of receiver **412** tracks the transmit clock of transmitter **410** with a phase delay. In other words, in principle, the receive clock of receiver **412** has the same frequency as the transmit clock of transmitter **410**, but with a fixed phase delay.

However, in the presence of jitter or a change in the cable response, these two clocks may have different instantaneous frequencies (frequency is derivative of phase with respect to time). This is due to the fact that, at the master **401**, the receiver **412** needs to dynamically change the relative phase of its receive clock with respect to the transmit clock of transmitter **410** in order to track jitter in the incoming signal from the transmitter **420** or to compensate for the change in cable response. Thus, in practice, the transmit and receive

clocks of the master **401** may be actually independent. At the master, this independence creates an asynchronous boundary between the transmit clock domain and the receive clock domain. By “transmit clock domain”, it is meant the region where circuit blocks are operated in accordance with transitions in the transmit clock signal TCLK. By “receive clock domain”, it is meant the region where circuit blocks are operated in accordance with transitions in the receive clock signal RCLK. In order to avoid any loss of data when data cross the asynchronous boundary between the transmit clock domain and the receive clock domain inside the master **401**, FIFOs are used at this asynchronous boundary. For the exemplary structure of the gigabit transceiver shown in FIG. 2, FIFOs **234** (FIG. 2) are placed at this asynchronous boundary. Since a constituent transceiver **108** (FIG. 1) is constructed such that it can be configured as a master or a slave, the FIFOs **234** (FIG. 2) are also included in the slave **402** (FIG. 4).

At the slave **402**, the transmit clock TCLK of transmitter **420** is phase locked to the receive clock RCLK of receiver **422**. Thus, TCLK may be different from GTX_CLK, a FIFO **430** is needed for proper transfer of data TXD from the MAC (not shown) to the transmitter **420**. The depth of the FIFO **430** must be sufficient to absorb any loss during the length of a data packet. The multiplexer **432** allows to use either the GTX_CLK or the receive clock RCLK of receiver **422** as the signal RX_CLK **429**. When the GTX_CLK is used as the RX_CLK **429**, the FIFO **434** is needed to ensure proper transfer of data RXD **427** from the receiver **422** to the MAC.

For the conceptual block diagram of FIG. 4, there are one transmit clock TCLK and one receive clock RCLK for a gigabit transceiver. The transmit clock TCLK is common to all four constituent transceivers since data transmitted simultaneously on all four twisted pairs of cable correspond to 4D symbols. Since data received from the four twisted pairs of cable are to be decoded simultaneously into 4D symbols, it is an efficient design to have all the digital processing blocks clocked by one clock signal RCLK. However, due the different cable responses of the four twisted pairs of cable, the A/D converter **216** (FIG. 2) of each of the four constituent transceivers requires a distinct sampling clock signal. Thus, in addition to the signals TCLK and RCLK, the gigabit transceiver system **200** requires four sampling clock signals.

There is an alternative structure for the gigabit transceiver where the partition of clock domains is different than the one shown in FIG. 2. This alternative structure (not shown explicitly) is similar to the one shown in FIG. 2 and only differs in that its transmit clock domain includes both the transmit clock domain and the receive clock domain of FIG. 2, and that the FIFO block **234** is not needed. In other words, in this alternative structure, the receive clock RCLK is the same as the transmit clock TCLK, and the transmit clock TCLK is used to clock both the transmitter and most of the receiver. The advantage of this alternative structure is that there is no asynchronous boundary between the transmit region and most of the receive region, thus allowing the echo canceller **232** and NEXT cancellers **230** to work with only one clock signal. The disadvantage of this alternative structure is that there is a potential for a performance penalty at the master when the constituent transceivers are tracking jitter. As a result of tracking jitter, the relative phase of a sampling clock signal with respect to the transmit clock TCLK may vary dynamically. This could cause the A/D converter to sample at noisy instants where transistors in circuit blocks operating according to the clock signal TCLK

are switching. Thus, the alternative structure is not as good as the structure shown in FIG. 2, with respect to the switching noise problem.

FIG. 5 is a simplified block diagram of an embodiment of the timing recovery system constructed according to the present invention and applied to the gigabit transceiver architecture of FIG. 2. The timing recovery system 222 (FIGS. 2 and 6) generates the different clock signals for the exemplary gigabit transceiver shown in FIG. 2, namely, the sampling clock signals ACLK0, ACLK1, ACLK2, ACLK3, the receive clock signal RCLK, and the transmit clock signal TCLK.

The timing recovery system 222 includes a set of phase detectors 502, 512, 522, 532, a set of loop filters 506, 516, 526, 536, a set of numerically controlled oscillators (NCO) 508, 518, 528, 538 and a set of phase selectors 510, 520, 530, 540, 550, 560. The adders 504, 514, 524, 534 are shown for conceptual illustration purpose only. In practice, these adders are implemented within the respective phase detectors 502, 512, 522, 532. The RCLK Offset is used to adjust the phase of the receive clock signal RCLK in order to reduce the effects of switching noise on the sampling operations of the corresponding A/D converters 216 (FIG. 2). Three of the four signals ACLK0 Offset, ACLK1 Offset, ACLK2 Offset, ACLK3 Offset are used to slightly adjust the phases of the respective sampling clocks ACLK0 through ACLK4 in order to further reduce these effects of switching noise. The phase adjustments of the receive clock RCLK and the sampling clocks ACLK0-3 are not a necessary function of the timing recovery system 222. However, the method and system for generating these phase adjustment signals constitute another novel aspect of the present invention and will be described in detail later.

Each of the phase detectors 502, 512, 522, 532 receives the corresponding 1D component of the 4D slicer error 42 (FIGS. 2 and 3) and the corresponding 1D component of the 4D tentative decision 44 (FIGS. 2 and 3) from the decoder 38 (FIG. 2) to generate a corresponding phase error. The phase errors 0 through 3 are inputted to the loop filters 506, 516, 526, 536, respectively. The loop filters 506, 516, 526, 536 generate and output filtered phase errors to the NCOs 508, 518, 528, 538. The loop filters 506, 516, 526, 536 can be of any order. In one embodiment, the loop filters are of second order. The NCOs 508, 518, 528, 538 generate phase control signals from the filtered phase errors. The phase selectors 510, 520, 530, 540 receive corresponding phase control signals from the NCOs 508, 518, 528, 538, respectively. Each of the phase selectors 510, 520, 530, 540 selects one out of several phases of the multi-phase signal 570 based on the value of the corresponding phase control signal, and outputs the corresponding sampling clock signal. In one embodiment of the invention, the multi-phase signal has 64 phases.

The multi-phase signal 570 is generated by a clock generator 580. In the exemplary embodiment illustrated in FIG. 5, the clock generator 580 includes a crystal oscillator 582, a frequency multiplier 584 and an 8-phase ring oscillator 586. The crystal oscillator 582 produces a 25 MHz clock signal. The frequency multiplier 584 multiplies the frequency of the 25 MHz clock signal by 40 and produces a 1 GHz clock signal. From the 1 GHz clock signal, the 8-phase ring oscillator 586 produces the 8 GHz 64-phase signal 570.

The receive clock signal RCLK, which is used to clock all the circuit blocks in the receive clock domain (which include all the digital signal processing circuit blocks in FIG. 2), can

be generated independently of the sampling clock signals ACLK0 through ACLK3. However, for design efficiency, RCLK is chosen to be related to one of the sampling clock signals ACLK0 through ACLK3. For the exemplary embodiment illustrated in FIG. 5, the receive clock signal RCLK is related to the sampling clock signal ACLK0. The receive clock signal RCLK is generated by inputting the sum of the phase control signal outputted from the NCO 508 and the RCLK Offset via an adder 542 to the phase selector 550. Based on this sum, the phase selector 550 selects one of the 64 phases of the multi-phase signal 570 and outputs the receive clock signal RCLK. Thus, when the RCLK Offset is zero, the receive clock signal RCLK is the same as the sampling clock ACLK0.

As discussed previously in relation to FIG. 4, when the constituent transceiver is configured as the master, its transmit clock TCLK is practically independent of its receive clock RCLK. In FIG. 5, when the constituent transceiver is the master, the transmit clock signal TCLK is generated by inputting the signal TCLK Offset, generated by the PHY Control system of the gigabit transceiver, to the phase selector 560. Based on the TCLK Offset, the phase selector 560 selects one of the 64 phases of the multi-phase signal 570 and produces the transmit clock signal TCLK. When the constituent transceiver is the slave, the transmit clock signal TCLK is generated by inputting the sum of the output of the NCO 508 and the signal TCLK Offset, via the adder 542, to the phase selector 560. Based on this sum, the phase selector 560 selects one of the 64 phases of the multi-phase signal 570 and produces the transmit clock signal TCLK. Thus, at the slave, the transmit clock signal TCLK and the receive clock signal RCLK are phase-locked (as discussed previously in relation to FIG. 4). In one embodiment of the present invention, the TCLK Offset is set equal to zero.

It is important to note that, referring to FIG. 5, the function performed by the combination of an NCO (508, 518, 528, 538) followed by a phase selector (510, 520, 530, 540, 550, 560) can be implemented by analog circuitry. The analog circuitry can be described as depicted in FIG. 5A. Each of the filtered phase errors outputted from the loop filters (506, 516, 526, 536) would be inputted to a D/A converter 508A, 518A, 528A and 538A, respectively, to be converted to analog form. Each of the analog filtered phase errors would then be inputted to a voltage-controlled oscillator (VCO) 510A, 520A, 530A and 540A, respectively. The VCOs 510A, 520A, 530A and 540A would produce the clock signals. The VCOs 510A, 520A, 530A and 540A can be implemented with well-known analog techniques such as those using varactor diodes.

FIG. 6 is a block diagram illustrating a detailed implementation of the phase detectors 502, 512, 522, 532, the loop filters 506, 516, 526, 536, and the NCOs 508, 518, 528, 538 of FIG. 5.

It is important to note that the 4D path connecting the phase detectors 502, 512, 522, 532, the loop filters 506, 516, 526, 536, the NCOs 508, 518, 528, 538 and the phase selectors 510, 520, 530, 540 (FIG. 5) can be thought of as the 4D forward path of a phase locked loop whose 4D feedback path goes from, referring now to FIG. 2, the A/D converters 216 to the demodulator 226 then back to the timing recovery 222. The input to this phase locked loop is actually phase information embedded in the slicer error 42 and tentative decision 44, and the phase locked loop output is the phases of the sampling clock signals. This phase locked loop is digital but can be approximated by a continuous-time phase locked loop for practical design analysis purpose, as long as the sampling rate is much larger than the bandwidth of the

loop. The theoretical transfer function of a continuous-time second-order phase locked loop is:

$$\frac{\Phi(s)}{\Theta(s)} = \frac{K_L \cdot s + K_L \cdot K_1}{s^2 + K_L \cdot s + K_L \cdot K_1}$$

where the transfer function of the loop filter is:

$$L(s) = K_L \cdot \left(1 + \frac{K_1}{s}\right) = K_v \cdot K_d \cdot \left(1 + \frac{K_1}{s}\right)$$

where K_v is the gain of the voltage-controlled oscillator, K_d is the gain of the phase detector, $K_L = K_v \cdot K_d$ and K_1 is the gain of the integrator inside the loop filter. For the digital phase locked loop of the present invention, the gain parameters K_v and K_1 can be computed from the word lengths and scale factors used in implementing the NCO and the integrator of the loop filter. However, the gain of the phase detector K_d is more conveniently computed by simulation. The gain parameters are used for the design and analysis of the digital phase locked loop.

FIG. 6 shows a phase detector 610, a first filter 630, a second filter 650, an adder 660 and an NCO 670. The phase detector 610 is an exemplary embodiment of the phase detectors 502, 512, 522, 532 of FIG. 5. The combination of the first filter 630, the second filter 650 and the adder 660 is an exemplary embodiment of the loop filters 506, 516, 526, 536 of FIG. 5. The NCO 670 is an exemplary embodiment of the NCOs 508, 518, 528, 538 of FIG. 5.

In FIGS. 6 through 8, the numbers in the form “Sn.k” indicate the format of the data, where S denotes a signed number, “n” denotes the total number of bits and “k” denotes the number of bits after the decimal point.

The phase detector 610 includes a lattice structure having two delay elements 612, 618, two multipliers 614, 620 and an adder 622. The phase detector 610 receives as inputs the corresponding 1D component of the 4D slicer error 42 (FIGS. 2 and 3) and the corresponding 1D component of the 4D tentative decision 44 (FIGS. 2 and 3) from the trellis decoder 38 (FIGS. 2 and 3). For simplicity, in FIG. 6, these two 1D components are labeled as 42A and 44A, respectively. It is understood that, for the phase detector of each of the four constituent transceivers of the gigabit transceiver, a distinct 1D component of the slicer error 42 and a distinct 1D component of the tentative decision 44 are used as inputs. On the upper branch of the lattice structure, the slicer error 42 is delayed by one unit of time (here, one symbol period) via the delay element 612, then multiplied by the tentative decision 44A to produce a pre-cursor phase error 615. The pre-cursor phase error 615, when accumulated over time, represents the correlation between a past slicer error and a present tentative decision, thus indicates the sampling phase error with respect to the zero-crossing point at the start of the signal pulse (this zero-crossing point is part of the pre-cursor introduced by design to the signal pulse by the precursor filter 28 of the FFE 26 in FIG. 2). On the lower branch of the lattice structure, the tentative decision 44A is delayed by one unit of time via the delay element 618, then multiplied by the slicer error 42A to produce a post-cursor phase error 621. The post-cursor phase error 621, when accumulated over time, represents the correlation between a present slicer error and a past tentative decision, thus indicates the sampling phase error with respect to the level-crossing point in the tail end of the signal pulse. In one embodiment, this level-crossing point is determined by the first tap coefficient of the DFE 312 of FIG. 3. At the zero-crossing point at the start of the signal pulse, the slope

of the signal pulse is positive, while at the level-crossing point at the tail end of the signal pulse, the slope of the signal pulse is negative. Thus, the pre-cursor phase error 615 and the post-cursor phase error 621 must be combined with opposite signs in the adder 622. The combination of the pre-cursor 615 and post-cursor phase errors 621 produces the phase error associated with one of the sampling clock signals ACLK0–ACLK3. This is the phase error indicated as one of the phase errors 0 through 3 in FIG. 5.

The phase offset 602 is one of the sampling clock offset signals ACLK0 Offset through ACLK3 Offset in FIG. 5. The phase offset 602, when needed, is generated by the PHY Control system of the gigabit transceiver. The phase offset 602 is delayed by one unit of time then is added to the combination of the pre-cursor error 615 and post-cursor 621 via the adder 622 to produce an adjusted phase error. The adjusted phase error 623 is stored in the delay element 624 and outputted to the first filter 630 at the next clock transition. The delay element 624 is used to prevent the propagation delay of the adder 622 from concatenating with the propagation delay of the adder 632 in the first filter 630.

The first filter 630, termed “phase accumulator”, accumulates the phase error 625 outputted by the phase detector 610 over a period of time then outputs the accumulated result at the end of the period of time. In the exemplary embodiment shown in FIG. 6, this period of time is 16 symbol periods. The first filter 630 is an “accumulate-and-dump” filter which includes the adder 632, a delay element (i.e., register) 634, and a 16-units-of-time register 636. The register 626 outputs a lowpass filtered phase error 627 at the rate of one per period of the TRSAMPO 604 clock, that is, one every 16 symbol periods. When the register 626 outputs the lowpass filtered phase error 627, the register 634 is cleared and the accumulation of phase error 625 restarts. It is noted that, downstream from the register 626, circuits are clocked at one sixteenth of the symbol rate.

The filtered phase error 637 is inputted to a multiplier 640 where it is multiplied by a factor different than 1 when it is desired that the bandwidth of the phase locked loop be different than its normal value (which is determined by the design of the filter). In the exemplary embodiment depicted in FIG. 6, filtered phase error 637 is multiplied by the value 2 outputted from a multiplexer 642 when the select signal 606 indicates that the loop filter bandwidth must be larger than normal value. This occurs, for example, during startup of the gigabit transceiver. Similarly, although not shown in FIG. 6, when it is desired that the loop filter bandwidth be narrower than normal value, the filtered phase error 637 can be multiplied by a value less than 1.

The output 644 of the multiplier 640 is inputted to the second filter 650 which is an integrator and to the adder 660. The integrator 650 is an IIR filter having an adder 652 and a register 654, operating at one sixteenth of the symbol rate. The integrator 650 integrates the signal 644 (which is essentially the filtered phase error 637) to produce an integrated phase error 656. The purpose of the phase locked loop is to generate a resulting phase for a sampling clock signal such that the phase error is equal to zero. The purpose of the integrator 650 in the phase locked loop is to keep the phase error of the resulting phase equal to zero even when there is static frequency error. Without the integrator 650, the static frequency error would result in a static phase error which would be attenuated but not made exactly zero by the phase locked loop. With the integrator 650 in the phase locked loop, any static phase error would be integrated to produce a large growing input signal to the NCO 670, which would cause the phase locked loop to correct the static phase

error. The integrated phase error **656** is scaled by a scale factor via a multiplier **658**. This scale factor contributes to the determination of the gain of the integrator **650**. The scaled result **659** is added to the signal **644** via an adder **660**.

The output **662** of the adder **660** is inputted to the NCO **670**. The output **662** is scaled by a scale factor, e.g., 2^{-5} , via a multiplier **672**. The resulting scaled signal is recursively filtered by an IIR filter formed by an adder **674** and a register **676**. The IIR filter operates at one sixteenth of the symbol rate. The signal **678**, outputted every 16 symbol periods, is used as the phase control signal to one of the phase selectors **510, 520, 530, 540, 550, 560** (FIG. 5).

For the embodiment shown in FIG. 6, the gain parameters discussed above are as follows. K_n , the gain of the NCO, is 2^{-11} for normal bandwidth mode, 2^{-10} for high bandwidth mode. K_i , the gain of the integrator **650**, is equal to the product of the scaling of the integrator register **654** (2^{-8} in FIG. 6) and the ratio of the phase locked loop sampling rate to the symbol rate (2^{-4} in FIG. 6). For the word lengths and scaling indicated in FIG. 6, K_d is equal to 2^{-12} . The gain K_d of the phase detector **610** is computed by simulations and is equal to 2.2. These parameters are used to compute the theoretical transfer function of the phase locked loop (PLL) which is then compared with the PLL transfer function obtained by simulation. The match is near perfect, confirming the validity of the design parameters.

One embodiment of the system **600** of FIG. 6 further includes the external control signals PLLFRZ, PLLPVAL, PLLPRST, PLLFVAL, PLLFRST, PLLPRAMP, which are not shown explicitly in FIG. 6.

The control signal PLLFRZ, when applied, forces the phase error to zero at point 1 of the first filter **630**, therefore causes freezing of updates of the frequency change and/or phase change, except for any phase change caused by a non-zero value in the frequency register **654** of the integrator **650**.

The control signal PLLPVAL is a 3-bit signal provided by the PHY Control system. It is used to specify the reset value of the NCO register **676** of the NCO **670**, and is used in conjunction with the control signal PLLPRST.

The control signal PLLPRST, when applied to the NCO register **676** in conjunction with the signal PLLPVAL, resets the 6 most significant bits of the NCO register **676** to a value specified by 8 times PLLPVAL. The reset is performed by stepping up or down the 6 MSB field of the NCO register **676** such that the specified value is reached after a minimum number of steps. Details of the phase reset logic block used to reset the value of the register **676** of the NCO **670** are shown in FIG. 7 and will be discussed later.

PLLFVAL is a 3-bit signal provided by the PHY Control system. It is to be interpreted as a 3-bit two's complement signed integer in the range $[-4,3]$. It is used to specify the reset value of the frequency register **654** of the integrator **650** and is used in conjunction with the control signal PLLFRST.

The control signal PLLFRST, when applied to the frequency register **654** of the integrator **650** in conjunction with the signal PLLFVAL, resets the frequency register **654** to the value 65536 times PLLFVAL.

The control signal PLLPRAMP loads the fixed number -2048 into the frequency register **654** of the integrator **650**. This causes the phase of a sampling clock signal (and receive clock RCLK) to ramp at the fixed rate of -2 ppm. This is used during startup at the master constituent transceiver. PLLPRAMP overrides PLLFRST. In other words, if both PLLPRAMP and PLLFRST are both applied, the value loaded into the frequency register **654** is -2048 , regardless of the value that PLLFRST tries to load.

FIG. 7 is a block diagram illustrating the phase reset logic block **700** to the NCO **670**. The control signal PLLPRST is applied to the AND gate **702**. The output of the AND gate **702** is applied to the increment/decrement enable input of the register **676**. The 3-bit value PLLPVAL from the PHY Control System of the gigabit transceiver is shifted left by 3 bits to form a 6-bit value **704**. The current output of the register **676** of the NCO **670** (FIG. 6), which is the phase control signal inputted to the corresponding phase selector (FIG. 5), is subtracted from this shifted value of PLLPVAL via an adder **706**. Module **708** determines whether the output of adder **706** is non-zero. If it is non-zero, then module **708** outputs a "1" to the AND gate **702** to enable the enable input of register **676**. If it is zero, module **706** outputs a zero to the AND gate **708** to disable the enable input of the register **676**. Module **710** determines whether the output of adder **706** is positive or negative. If it is positive, module **710** outputs a count up indicator to the register **676**. If it is negative, module **710** outputs a count down indicator to register **676**.

The subtraction at adder **706** finds the shortest path from the current value of the NCO register **676** to the shifted PPLVAL **704**. For example, suppose the current phase value of register **676** is 20. If the shifted PPLVAL **704** (which is the desired value) is 32, the difference is 12, which is positive, therefore, the register **676** is incremented. If the desired phase value is 56, the difference is 36 or "100100" which is interpreted as -28 , so the register **676** will be decremented 28 consecutive times. The phase steps occur at the rate of one every 16 symbol periods. This single stepping is needed because of the way the phase selector operates. The phase selector can only increment or decrement from its current setting.

FIG. 8 is a block diagram of an exemplary phase shifter logic block used for the phase control of the receive clock signal RCLK. The phase shifter logic block **800** is needed when the signal RCLK Offset (FIG. 5) is used to adjust the phase of the receive clock signal RCLK. The signal RCLK Offset is a 6-bit signal provided by the PHY Control system, and specifies the amount by which the phase of RCLK must be shifted. Even if the signal RCLK Offset indicates a large amount of phase shift, this phase shift must be transferred to the input of the phase selector **550** (FIG. 5) one step at a time due to the way the phase selector operates. The change of phase of RCLK must occur in the direction indicated by a control signal STEPDIR generated by the PHY Control system.

The phase shifter logic block **800** includes a comparator **802**, an offset register **804** and the adder **542** (the same adder indicated in FIG. 5). The comparator **802** compares the output **806** of the offset register **804** with the signal RCLK Offset. If the two signals are equal, then the comparator **802** outputs a "0" to the enable input of the offset register **804** to disable the up/down counting of the offset register **804**, thus keeping the output **806** the same for the next time period. If the two signals are not equal, the comparator **802** outputs a "1" to the enable input of the offset register **804** to enable the up/down counting, causing the output **806** to be incremented or decremented at the next time period. The signal STEPDIR from the PHY Control system is inputted to the up/down input of the offset register **804** to control the counting direction. The output **806** from the offset register **804** is added to the phase control signal **509** produced by the NCO **508** (FIG. 5) via the adder **542** to generate the phase control signal **549** (FIGS. 8 and 5) for the RCLK phase selector **550** (FIG. 5).

The coupling of switching noise from the digital signal processor that implements the transceiver functions to each

of the A/D converters is an important problem that needs to be addressed. Switching noise occurs when transistors switch states in accordance with transitions in the clock signal (or signals) that controls their operation. Switching noise in the digital section of the transceiver can be coupled to the analog section of the transceiver. Switching noise can cause severe degradation to the performance of an A/D converter if it occurs right at or near the instant the A/D converter is sampling the received signal. The present invention, in addition to providing a timing recovery method and system, also provides a method and system for minimizing the degradation of the performance of the A/D converters caused by switching noise.

The effect of switching noise on an A/D converter can be reduced if the switching noise is synchronous (with a phase delay) with the sampling clock of the A/D converter. If, in addition, it is possible to adjust the phase of the sampling clock of the A/D converter with respect to the phase of the switching noise, then the phase of the sampling clock of the A/D converter can be optimized for minimum noise. It is noted that, for a local gigabit transceiver, the sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3** are synchronous to each other (i.e., having the same frequency) because they are synchronous to the 4 transmitters of the remote transceiver and these 4 remote transmitters are clocked by a same transmit clock signal **TCLK**. It is also important to note that the local receive clock signal **RCLK** is synchronous to the local sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3**.

Referring to FIGS. 2 and 5, the four A/D converters **216** of the four constituent transceivers are sampled with the sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3**. Each of the phases of these sampling clock signals is determined by the subsystem **600** (FIG. 6) of the timing recovery system **222** in response to the phase of the corresponding received signal, which depends on the remote transmitter and the line characteristics. Thus, the phases of the sampling clock signals change from line to line, and are not under the control of the system designer.

However, the relative phase of the receive clock signal **RCLK** with respect to the sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3** can be controlled by adjusting the signal **RCLK Offset** (FIG. 5). The signal **RCLK Offset** can be used to select the **RCLK** phase that would cause the least noise coupling to the A/D converters **216** of FIG. 2. The underlying principle is the following. Referring to FIG. 2 and the boundaries of the clock domain, the entire digital signal processing, control and interface functions of the receiver operate in accordance with transitions in the receive clock signal **RCLK**. In other words, most of the digital logic circuits switch states on a transition of **RCLK** (more specifically, on a rising edge of **RCLK**). Only a small portion of the transceiver operates in accordance with transitions in the transmit clock signal **TCLK**. Therefore, most of the switching noise is synchronous with the receive clock signal **RCLK**. Since the receive clock signal **RCLK** is synchronous with the sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3**, it follows that most of the switching noise is synchronous with the sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3**. Therefore, if the phase of the receive clock signal **RCLK** is adjusted such that a transition in the signal **RCLK** occurs as far as possible in time from each of the sampling clock signals **ACLK0**, **ACLK1**, **ACLK2**, **ACLK3**, then the switching noise coupling to the A/D converters will be minimized.

The process for adjusting the phase of the receive clock signal **RCLK** can be summarized as follows. The process

performs an exhaustive search over all the **RCLK** phases that, by design, can possibly exist in one symbol period. For each phase, the process computes the sum of the mean squared errors (MSEs) of the 4 pairs (i.e., the 4 constituent transceivers). At the end of the search, the process selects the **RCLK** phase that minimizes the sum of the MSEs of the four pairs. The following is a description of one embodiment of the **RCLK** phase adjustment process, where there are 64 possible **RCLK** phases.

FIG. 9 is a flowchart illustrating the process **900** for adjusting the phase of the receive clock signal **RCLK**. Upon Start (block **902**), process **900** initializes all the state variables (which include counters, registers), sets **Offset** to -32 (block **904**), sets **Min_MSE** equal to the MSE of the gigabit transceiver before any **RCLK** phase change, and sets **BestOffset** equal to zero. The MSE of the gigabit transceiver is the sum of the mean squared errors (MSEs) of the 4 constituent transceivers. The MSE of a constituent transceiver is the mean squared error of the corresponding 1D component of the 4D slicer error **42** (FIG. 2), and is outputted by a MSE computation block **1200** (FIG. 12) for every frame. Each frame is equal to 1024 symbol periods. This initialization is done within a duration of 1 frame. Process **900** then waits for the effect of the **RCLK** phase change on the system to settle (block **906**). The duration of this waiting is 5 frames. Process **900** then computes MSE (by summing the MSEs of all four constituent transceivers outputted by the corresponding MSE computation block **1200** of FIG. 12) which corresponds to the current setting of **RCLK Offset** (block **908**). The duration of block **908** is one frame. In block **910**, process **900** compares the new MSE with **Min_MSE**. If the new MSE is strictly less than **Min_MSE**, then **Min_MSE** is set to the value of the new MSE and **BestOffset** is set to the value of **Offset**. In block **912**, process checks whether **Offset** is equal to 31, i.e., whether all possible 64 phase offsets have been searched. If **Offset** is not equal to 31, then process **900** increments **Offset** by 1 (block **914**) then continues the search for the best **RCLK Offset** by going back to block **906**. If **Offset** is equal to 31, that is, if process **900** has searched all possible 64 phase offsets, then process **900** sets **Offset** equal to the value of **BestOffset** (block **916**) then terminates (block **918**). The duration of each of blocks **914** and **916** is 1 frame.

After adjustment of the receive clock **RCLK** phase, small adjustments can be made to the phases of the sampling clocks **ACLK1**, **ACLK2**, **ACLK3** to further reduce the coupling of switching noise to the A/D converters. Since the timing recovery system **222** of FIG. 5 without the **ACLK0-3** Offsets, through the phase locked loop principle, already sets the sampling clocks at the optimal sampling positions with respect to the pulse shape of incoming signals from the remote transceivers, the small phase adjustments made to the sampling clocks could cause some loss of performance of the A/D converters. However, the net result is still better than performing no phase adjustment of the sampling clocks and allowing the A/D converters to sample the incoming signals at a noisy instant where the transistors in the digital section are switching states. In the embodiment depicted in FIG. 5, phase adjustment is not made to the sampling clock **ACLK0** because, by design of the structure of the embodiment, the phase difference between **ACLK0** and **RCLK** is equal to **RCLK Offset**. Thus, in this embodiment, any adjustment to the phase of **ACLK0** will also move **RCLK** away from the optimal position determined by process **900** above by the same amount of phase adjustment.

FIGS. 10A, 10B, 10C illustrate three examples of distribution of the transitions of clock signals within a symbol

period to further clarify the concept of phase adjustment of the clock signals. It is noted that, in these examples, the four sampling clock signals **ACLK0–3** are shown as occurring in their consecutive order within a symbol period for illustrative purpose only. It is understood that the sampling clock signals **ACLK0–3** can occur in any order.

FIG. 10A is a first example of clock distribution where the transitions of the four sampling clock signals **ACLK0–3** are evenly distributed within the symbol period of 8 nanoseconds (ns). Thus, each **ACLK** clock transition is 2 ns apart from an adjacent transition of another **ACLK** clock. Therefore, for this clock distribution example, a transition of the receive clock **RCLK** can only be placed at most 1 ns away from an adjacent **ACLK** transition. This “distance” (phase delay) may not be enough to reduce the coupling of switching noise to the two A/D converters associated with the two adjacent sampling clock signals (**ACLK3** and **ACLK0**, in the example). In this case, it may be desirable to slightly adjust the phase of the two adjacent sampling clock signals to move their respective transitions further away from a **RCLK** transition, as illustrated by their new transition occurrences within a symbol period in **FIG. 10A**.

FIG. 10B is a second example of clock distribution where the transitions of the four sampling clock signals **ACLK0–3** are distributed within the symbol period of 8 nanoseconds (ns) such that each **ACLK** clock transition is 1 ns apart from an adjacent transition of another **ACLK** clock. For this clock distribution example, a transition of the receive clock **RCLK** can be positioned midway between the last **ACLK** transition of one symbol period (**ACLK3** in **FIG. 10B**) and the first **ACLK** transition of the next symbol period (**ACLK0** in **FIG. 10B**) so that the **RCLK** transition is 2.5 ns from an adjacent **ACLK** transition. This “distance” (phase delay) may be enough to reduce the coupling of switching noise to the two A/D converters associated with the two adjacent sampling clock signals (**ACLK3** and **ACLK0**, in the example). In this case, phase adjustment of the two adjacent sampling clock signals to move their respective transitions further away from a **RCLK** transition may not be needed.

FIG. 10C is a third example of clock distribution where the transitions of the four sampling clock signals **ACLK0–3** occur at the same instant within the symbol period of 8 nanoseconds (ns). In this clock distribution example, a transition of the receive clock **RCLK** can be positioned at the maximum possible distance of 4 ns from an adjacent **ACLK** transition. This is the best clock distribution that allows maximum reduction of coupling of switching noise to the four A/D converters associated with the sampling clock signals. In this case, there is no need for phase adjustment of the sampling clock signals.

For the embodiment shown in **FIG. 5** of the timing recovery system **222** (**FIG. 2**), the following phase adjustment process is applied to the three sampling clock signals **ACLK1**, **ACLK2**, **ACLK3**. It is understood that, in a different embodiment of the timing recovery system **222** (**FIG. 2**) where the receive clock signal **RCLK** is not tied to one of the sampling clock signals **ACLK0–3**, the following phase adjustment process can be applied to all of the sampling clock signals.

The process for adjusting the phase of a sampling clock signal **ACLKx** (“x” in **ACLKx** denotes one of **0,1,2,3**) can be summarized as follows. The process performs a search over a small range of phases around the initial **ACLKx** phase. For each phase, the process logs the mean squared error **MSE** of the associated constituent transceivers. At the end of the search, the process selects the **ACLKx** phase that minimizes the **MSE** of the associated constituent transceiver.

Whenever the phase of a sampling clock signal **ACLKx** changes, the coefficients of the echo canceller **232** and of the **NEXT** cancellers **230** change. Thus, to avoid degradation of performance, the phase steps of the sampling clocks should be small so that the change they induce on the coefficients is also small. When the phase adjustment requires multiple consecutive phase steps, the convergence of the coefficients of the echo canceller **232** and of the **NEXT** cancellers **230** should be fast in order to avoid a buildup of coefficient mismatch.

FIG. 11 is a flowchart illustrating an embodiment of the process for adjusting the phase of a sampling clock signal **ACLKx** associated with one of the constituent transceivers, where the search is over a range of 16 phases around the initial **ACLKx** phase. For each of the constituent transceivers, process **1100** of **FIG. 11** is run independently of and concurrently with the other constituent transceivers. Upon Start (block **1102**), process **1100** initializes all the state variables (which include counters, registers), sets **Offset** to **-8** (block **1104**), sets **Min_MSE** equal to the **MSE** of the associated constituent transceiver before any **RCLK** phase change, and sets **BestOffset** equal to zero. The **MSE** of the associated constituent transceiver is the mean squared error of the corresponding 1D component of the 4D slicer error **42** (**FIG. 2**). This initialization is done within a duration of 1 frame. Process **1100** then waits for the effect of the **ACLK** phase change on the system to settle (block **1106**). The duration of this waiting is 32 frames. (block **1108**). The duration of block **1108** is one frame. In block **1110**, process **1100** compares the new **MSE** (outputted by the corresponding **MSE** computation block **1200** of **FIG. 12**) which corresponds to the current setting of **ACLKx** **Offset** with **Min_MSE**. If the new **MSE** is strictly less than **Min_MSE**, then **Min_MSE** is set to the value of the new **MSE** and **BestOffset** is set to the value of **Offset**. In block **1112**, process **1100** checks whether **Offset** is equal to 7, i.e., whether all 16 phase offsets in the range have been searched. If **Offset** is not equal to 7, then process **1200** increments **Offset** by 1 (block **1114**) then continues the search for the best **ACLKx** **Offset** by looping back to block **1106**. If **Offset** is equal to 7, that is, if process **1100** has searched all the 16 phase offsets in the range, then process **1100** sets **Offset** equal to the value of **BestOffset** (block **1116**) then terminates (block **1118**). The duration of each of blocks **1114** and **1116** is 1 frame.

FIG. 12 is a block diagram of an exemplary implementation of the **MSE** computation block used for computing the mean squared error of a constituent transceiver. In one embodiment of the gigabit transceiver, there are four **MSE** computation blocks, one for each of the four constituent transceivers. The four **MSE** computation blocks are run independently and concurrently for the four constituent transceivers. The **MSE** computation block **1200** includes a squaring module **1202** and an infinite impulse response (**IIR**) filter **1204**. The **IIR** filter **1204** includes an adder **1206**, a feedback delay element **1208** and a forward delay element **1210**. The squaring module **1202** receives the corresponding 1D component of the 4D slicer error **42** (**FIG. 2**), which is denoted as **42A** for simplicity, and outputs the squared error value to the filter **1204**. The filter **1204** accumulates the squared error values by adding via the adder **1206** the current squared error value to the previous squared error value stored in the feedback delay element **1208**. The accumulated value is stored in the forward register **1210**. In the exemplary embodiment shown in **FIG. 12**, the squared error values are accumulated for **1024** symbol periods (which is one frame of the **PHY** Control system). Since the

accumulation period is sufficiently long, the accumulated value practically corresponds to the mean squared error. At the end of the accumulation period, the clock signal 1220 from the PHY Control system clears the contents of the feedback delay element, and clocks the forward delay element 1210 so that the forward delay element 1210 outputs the accumulated value MSE and resets to zero.

While certain exemplary embodiments have been described in detail and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention. It will thus be recognized that various modifications may be made to the illustrated and other embodiments of the invention described above, without departing from the broad inventive scope thereof. It will be understood, therefore, that the invention is not limited to the particular embodiments or arrangements disclosed, but is rather intended to cover any changes, adaptations or modifications which are within the scope and spirit of the invention as defined by the appended claims.

What is claimed is:

1. A method for generating a set of clock signals in a system, the set of clock signals comprising a set of sampling clock signals, the system comprising a set of subsystems, each of the subsystems comprising an analog section, each of the analog sections operating in accordance with a corresponding one of the sampling clock signals, the method comprising:

- (a) generating a phase error for each of the sampling clock signals from a corresponding phase detector;
- (b) inputting each of the phase errors to a corresponding loop filter;
- (c) generating filtered phase errors from the corresponding loop filters;
- (d) inputting each of the filtered phase errors to a corresponding oscillator;
- (e) generating phase control signals from the corresponding oscillators;
- (f) inputting each of the phase control signals to a corresponding phase selector;
- (g) generating the sampling clock signals from the corresponding phase selectors; and
- (h) passing the sampling clock signals to the respective subsystems.

2. The method of claim 1 wherein the set of clock signals further comprises a receive clock signal and wherein each of the subsystems further comprises a digital section, the digital sections operating in accordance with the receive clock signal.

3. The method of claim 2 wherein the receive clock signal is related to one of the sampling clock signals.

4. The method of claim 3 further comprising:

- (1) combining one of the phase control signals with a receive clock offset to produce a phase shift value;
- (2) inputting the phase shift value to a receive clock phase selector; and
- (3) generating the receive clock signal from the receive clock phase selector.

5. The method of claim 4 wherein the phase shift value comprises a set of phase steps and wherein operation (2) comprises inputting one phase step of the phase shift value at a time to the receive clock phase selector.

6. The method of claim 1 wherein the set of clock signals further comprises a transmit clock signal and wherein each of the subsystems further comprises a transmit section, the transmit sections operating in accordance with the transmit clock signal.

7. The method of claim 6 further comprising:
inputting a transmit clock offset to a transmit clock phase selector; and

generating the transmit clock signal from the transmit clock phase selector.

8. The method of claim 7 wherein the transmit clock offset is equal to zero.

9. The method of claim 6 wherein the transmit clock signal is related to one of the sampling clock signals.

10. The method of claim 9 further comprising:

inputting one of the phase control signals to a transmit clock phase selector; and

generating the transmit clock signal from the transmit clock phase selector.

11. The method of claim 1 wherein each of the phase detectors receives a corresponding slicer error and a corresponding tentative decision from a decoding system.

12. The method of claim 11 wherein operation (a) comprises:

(1) generating a pre-cursor phase error by multiplying the corresponding tentative decision by a delayed version of the corresponding slicer error;

(2) generating a post-cursor phase error by multiplying the corresponding slicer error by a delayed version of the corresponding tentative decision;

(3) combining the pre-cursor and post-cursor phase errors to produce the corresponding phase error.

13. The method of claim 12 wherein operations (1), (2) and (3) are performed via a lattice structure, the lattice structure comprising two delay elements, two multipliers and an adder.

14. The method of claim 12 wherein operation (3) comprises combining the pre-cursor, post-cursor phase errors and an offset input from a control unit to produce the corresponding phase error.

15. The method of claim 1 wherein operation (c) comprises:

(1) accumulating a number of consecutive values of one of the phase errors via a first filter, resulting in a sum value;

(2) outputting the sum value from the first filter;

(3) integrating the sum value via a second filter to produce an integral value; and

(4) combining the sum value and the integral value to produce a filtered phase error.

16. The method of claim 15 wherein operation (3) comprises scaling the integrated sum value by a scale factor to produce the integral value.

17. The method of claim 15 wherein operation (c) further comprises, before operation (3), multiplying the sum value by a factor different than 1 when the system is operating in a different bandwidth mode.

18. The method of claim 1 wherein operation (e) comprises filtering recursively the filtered phase errors to produce the corresponding phase control signals.

19. The method of claim 18 wherein operation (e) further comprises scaling, before filtering recursively, the filtered phase errors by a scale factor.

20. The method of claim 1 wherein operation (g) comprises:

(1) inputting a multi-phase input signal from a clock generator to each of the phase selectors; and

(2) selecting at each of the phase selectors one of the phases of the multi-phase input signal based on the phase control signal received from the corresponding oscillator.

21. A timing recovery system for generating a set of clock signals in a processing system, the set of clock signals comprising a set of sampling clock signals, the processing system comprising a set of processing subsystems, each of the processing subsystems comprising an analog section, each of the analog sections operating in accordance with a corresponding one of the sampling clock signals, the timing recovery system comprising:

- (a) a set of phase detectors generating phase errors for the corresponding sampling clock signals;
- (b) a set of loop filters coupled to the corresponding phase detectors, the loop filters receiving the corresponding phase errors and generating filtered phase errors;
- (c) a set of oscillators coupled to the corresponding loop filters, the oscillators receiving the filtered phase errors and generating phase control signals; and
- (d) a set of phase selectors coupled to the corresponding oscillators, the phase selectors receiving the phase control signals and generating the sampling clock signals, the phase selectors being connected to the respective subsystems to deliver the sampling clock signals to the respective subsystems.

22. The timing recovery system of claim 21 wherein the set of clock signals further comprises a receive clock signal and wherein each of the processing subsystems further comprises a digital section, the digital sections operating in accordance with the receive clock signal.

23. The timing recovery system of claim 22 wherein the receive clock signal is related to one of the sampling clock signals.

24. The timing recovery system of claim 23 further comprising a first adder and a receive clock phase selector, the first adder receiving one of the phase control signals and a receive clock offset and generating a phase shift value, the receive clock phase selector receiving the phase shift value and generating the receive clock signal.

25. The timing recovery system of claim 24 wherein the phase shift value comprises a set of phase steps and wherein the receive clock phase selector receives the phase shift value in the form of consecutive phase steps.

26. The timing recovery system of claim 21 wherein the set of clock signals further comprises a transmit clock signal and wherein each of the subsystems further comprises a transmit section, the transmit sections operating in accordance with the transmit clock signal.

27. The timing recovery system of claim 26 further comprising a transmit clock phase selector, the transmit clock phase selector receiving a transmit clock offset and generating the transmit clock signal.

28. The timing recovery system of claim 27 wherein the transmit clock offset is equal to zero.

29. The timing recovery system of claim 26 wherein the transmit clock signal is related to one of the sampling clock signals.

30. The timing recovery system of claim 29 further comprising a transmit clock phase selector, the transmit

clock phase selector receiving one of the phase control signals and generating the transmit clock signal.

31. The timing recovery system of claim 21 wherein each of the phase detectors receives a corresponding slicer error and a corresponding tentative decision from a decoding system.

32. The timing recovery system of claim 31 wherein each of the phase detectors comprises a lattice structure, the lattice structure comprising two delay elements, two multipliers and an adder, the lattice structure generating a pre-cursor phase error by multiplying the corresponding tentative decision by a delayed version of the corresponding slicer error and generating a post-cursor phase error by multiplying the corresponding slicer error by a delayed version of the corresponding tentative decision and combining the pre-cursor and post-cursor phase errors to produce the corresponding phase error.

33. The timing recovery system of claim 32 wherein at least one of the phase detectors further receives an offset input from a control unit and wherein the associated lattice structure combines the pre-cursor, post-cursor phase errors and the offset input to produce the corresponding phase error.

34. The timing recovery system of claim 21 wherein at least one of the loop filters comprises a first filter for accumulating a number of consecutive values of one of the phase errors to produce a filtered phase error.

35. The timing recovery system of claim 21 wherein at least one of the loop filters comprises a first filter for accumulating a number of consecutive values of one of the phase errors to produce a sum value, a second filter for integrating the sum value to produce an integral value and an adder for combining the sum value and the integral value to produce a filtered phase error.

36. The timing recovery system of claim 35 wherein the second filter includes a multiplier for scaling the integrated sum value by a scale factor to produce the integral value.

37. The timing recovery system of claim 35 wherein at least one of the loop filters further comprises a multiplier for multiplying the sum value by a factor different than 1 when the system is operating in a different bandwidth mode.

38. The timing recovery system of claim 21 wherein each of the oscillators comprises an infinite impulse response filter for filtering recursively the filtered phase errors to produce the corresponding phase control signals.

39. The timing recovery system of claim 38 wherein at least one of the oscillators further comprises a multiplier for scaling the filtered phase errors by a scale factor and outputting the scaled filtered phase errors to the associated impulse response filter.

40. The timing recovery system of claim 21 wherein each of the phase selectors receives a multi-phase input signal from a clock generator and selects one of the phases of the multi-phase input signal based on the phase control signal received from the corresponding oscillator.